

NBS
PUBLICATIONS

A11102 615982

NATL INST OF STANDARDS & TECH R.I.C.



A11102615982

/Theory and application of expert system
QC100 .U57 NO.717 V1986 C.2 C.1 NBS-PUB-



NBS SPECIAL PUBLICATION **717**

U.S. DEPARTMENT OF COMMERCE/National Bureau of Standards

Theory and Application of Expert Systems in Emergency Management Operations

Gass and Chapman, Editors



QC
100
.U57
NO.717
1986
C.2

T

he National Bureau of Standards¹ was established by an act of Congress on March 3, 1901. The Bureau's overall goal is to strengthen and advance the nation's science and technology and facilitate their effective application for public benefit. To this end, the Bureau conducts research and provides: (1) a basis for the nation's physical measurement system, (2) scientific and technological services for industry and government, (3) a technical basis for equity in trade, and (4) technical services to promote public safety. The Bureau's technical work is performed by the National Measurement Laboratory, the National Engineering Laboratory, the Institute for Computer Sciences and Technology, and the Institute for Materials Science and Engineering.

The National Measurement Laboratory

Provides the national system of physical and chemical measurement; coordinates the system with measurement systems of other nations and furnishes essential services leading to accurate and uniform physical and chemical measurement throughout the Nation's scientific community, industry, and commerce; provides advisory and research services to other Government agencies; conducts physical and chemical research; develops, produces, and distributes Standard Reference Materials; and provides calibration services. The Laboratory consists of the following centers:

- Basic Standards²
- Radiation Research
- Chemical Physics
- Analytical Chemistry

The National Engineering Laboratory

Provides technology and technical services to the public and private sectors to address national needs and to solve national problems; conducts research in engineering and applied science in support of these efforts; builds and maintains competence in the necessary disciplines required to carry out this research and technical service; develops engineering data and measurement capabilities; provides engineering measurement traceability services; develops test methods and proposes engineering standards and code changes; develops and proposes new engineering practices; and develops and improves mechanisms to transfer results of its research to the ultimate user. The Laboratory consists of the following centers:

- Applied Mathematics
- Electronics and Electrical Engineering²
- Manufacturing Engineering
- Building Technology
- Fire Research
- Chemical Engineering²

The Institute for Computer Sciences and Technology

Conducts research and provides scientific and technical services to aid Federal agencies in the selection, acquisition, application, and use of computer technology to improve effectiveness and economy in Government operations in accordance with Public Law 89-306 (40 U.S.C. 759), relevant Executive Orders, and other directives; carries out this mission by managing the Federal Information Processing Standards Program, developing Federal ADP standards guidelines, and managing Federal participation in ADP voluntary standardization activities; provides scientific and technological advisory services and assistance to Federal agencies; and provides the technical foundation for computer-related policies of the Federal Government. The Institute consists of the following centers:

- Programming Science and Technology
- Computer Systems Engineering

The Institute for Materials Science and Engineering

Conducts research and provides measurements, data, standards, reference materials, quantitative understanding and other technical information fundamental to the processing, structure, properties and performance of materials; addresses the scientific basis for new advanced materials technologies; plans research around cross-country scientific themes such as nondestructive evaluation and phase diagram development; oversees Bureau-wide technical programs in nuclear reactor radiation research and nondestructive evaluation; and broadly disseminates generic technical information resulting from its programs. The Institute consists of the following Divisions:

- Ceramics
- Fracture and Deformation³
- Polymers
- Metallurgy
- Reactor Radiation

¹Headquarters and Laboratories at Gaithersburg, MD, unless otherwise noted; mailing address Gaithersburg, MD 20899.

²Some divisions within the center are located at Boulder, CO 80303.

³Located at Boulder, CO, with some elements at Gaithersburg, MD.

Theory and Application of Expert Systems in Emergency Management Operations

**Proceedings of a Symposium
Held at the Department of Commerce
Washington, DC
April 24-25, 1985**

NBS
RESEARCH
INFORMATION
CENTER

NBS
QC100
.L57
mb.717
1986
C.2

Edited by:

Saul I. Gass

College of Business and Management
University of Maryland
College Park, Maryland 20742

and

Robert E. Chapman

National Engineering Laboratory
National Bureau of Standards
Gaithersburg, Maryland 20899

Sponsored by:

Federal Emergency Management Agency
500 C Street, S.W.
Washington, DC 20472

and

Center for Applied Mathematics
National Engineering Laboratory
National Bureau of Standards
Gaithersburg, Maryland 20899



U.S. DEPARTMENT OF COMMERCE, Malcolm Baldrige, Secretary
NATIONAL BUREAU OF STANDARDS, Ernest Ambler, Director

Issued November 1986

Library of Congress Catalog Card Number: 86-600546
National Bureau of Standards Special Publication 717
Natl. Bur. Stand. (U.S.), Spec. Publ. 717, 264 pages (Nov. 1986)
CODEN: XNBSAV

U.S. GOVERNMENT PRINTING OFFICE
WASHINGTON: 1986



Federal Emergency Management Agency

Washington, D.C. 20472

Foreword

by

The Honorable Julius W. Becton, Jr.

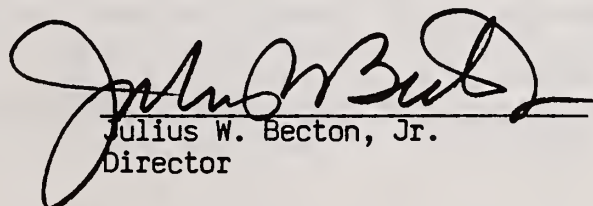
Director

Federal Emergency Management Agency

As the recently appointed director of the Federal Emergency Management Agency (FEMA), I would like to reinforce the thought-provoking material contained in these Proceedings of the Symposium on the Theory and Application of Expert Systems in Emergency Management Operations. As pointed out in the symposium, expert systems hold unique advantages in the arena of emergency management. These systems can give you advice quicker, better and more extensively than panicked human beings. Quicker, because they can be consulted at any hour by using a personal computer. Better, because they contain the distilled opinion of a wide range of expertise. More extensively, because the computer can simultaneously consider a wider range of input factors than any human being. And the computer never panics, seldom sleeps, and has never been known to be absent.

In the time that has passed since the symposium, FEMA has taken some of the first steps indicated by the symposium. We now have available on our computer or under development expert systems in such diverse areas as notification of emergency weather conditions and estimation of facility damage in earthquakes. The technology is being used to codify emergency expertise, and to improve our emergency management capabilities.

The future is limited only by our imagination. One can imagine the ability to call a number and to be answered by a computer that can understand your speech. This computer would have a suite of expert systems which it would relate to your description of your problem, and give you the best advice available. The ability to perform smarter with less is and will continue to be a driving force. The publication of this document is evidence of our continued commitment to excellent research devoted to fulfilling human needs.



Julius W. Becton, Jr.
Director

ABSTRACT


The First Symposium on The Theory and Application of Expert Systems in Emergency Management, held at the Department of Commerce, Washington, DC (April 24 and 25, 1985) was funded by the Federal Emergency Management Agency and organized by the National Bureau of Standards' Operations Research Division. The purpose of the symposium was to bring together researchers in expert systems, artificial intelligence, and emergency operations in a forum to review the concepts of expert systems and the problems of emergency management, with the objective of determining how expert systems can be used to augment the experience of local, State and Federal emergency managers faced with the difficult tasks of determining the best response to an emergency situation. Speakers addressed the following areas: The theory and uncertainty aspects of expert systems, artificial intelligence's future role in emergency management, technology for building and using expert systems, emergency management decisions and information needs and uses, applications of expert systems in the management of chemical spills and shipboard and coal mine fires, and the role and use of simulation in emergency management expert systems.

Keywords: Artificial intelligence; chemical spills; coal mine fires; decision support systems; emergency management; expert systems; fire management; shipboard fires; simulation; uncertainty

Note: The paper "Emergency Management for Chemical Spills" by C. K. Johnson that was presented at the Symposium is not included in this volume. The interested reader should consult "Emergency Management of Inland Oil and Hazardous Chemical Spills: A Case Study in Knowledge Engineering," Chapter 10 in Building Expert System, F. Hayes-Roth, D. A. Waterman and Douglas B. Lenat (eds.), Addison-Wesley Publishing Company, Inc., 1983.

TABLE OF CONTENTS

	Page
Foreword -- Julius W. Becton, Jr.....	iii
Abstract	iv
Welcome -- Louis O. Giuffrida.....	1
Introductory Remarks -- Raymond G. Kammer.....	4
A Theory of Expert Systems -- Andrew B. Whinston, Keith Hall and James Moore.....	6
Management of Uncertainty in Expert Systems -- Lotfi A. Zadeh.....	22
New Tools for Emergency Management -- David L. Waltz.....	40
TIMM - A Tool for Building Expert Systems -- Joseph N. Craig.....	57
Decision Support Systems and Expert Systems for Emergency Management Operations: A Microcomputer Approach -- Frank de Balogh.....	86
Optimum Emergency Management: The Effective Use of Information Technology -- Robert Lee Chartrand.....	109
Roles of Simulation in the Application of Expert Systems in Emergency Management Operations -- A. Ben Clymer.....	124
Expert System for Fire Emergency Management -- Geoffrey N. Berlin.....	158
The VENTILATION EXPERT for Underground Coal Mines -- Jeffrey L. Kohler..	183
Emergency Management Decisions at the Local Level -- Robert S. Wilkerson.....	203
The SMOKEY Project: An Intelligent Sensor-Based Fireman's Assistant -- Jaime G. Carbonell.....	211
Symposium Program.....	223
Speakers Names and Addresses.....	226
Appendix I: Expert Systems and Emergency Management: Background Paper -- Saul I. Gass and Robert E. Chapman.....	227
Appendix II: Expert Systems and Emergency Management: Research Directions -- Saul I. Gass and Robert E. Chapman.....	250



WELCOME

Honorable Louis O. Giuffrida
Director
Federal Emergency Management Agency

Good morning, ladies and gentlemen. I would like to welcome you to the first "Expert Systems in Emergency Management Symposium." One of the questions that you could ask is: "Why should the Federal Emergency Management Agency co-sponsor a symposium on expert systems now, especially during a time of budget constraints?" Because of a large and expanding deficit we're having to cut back funding in many areas. Yet, we are funding this symposium.

There are several good answers to this question. The first answer is that this technology is just beginning to emerge from the laboratory into practical application. We need to plan this type of information system now in order to obtain good performance later.

However, that is not the only reason for beginning the planning now. As I'm sure many of you know the common definition of an expert is an ordinary drip under pressure. However, it has been shown in field after field, from quarterbacking to fighter pilots to heart surgery that the true definition of experts are those 20 percent of the people who do fifty percent of the successful work. Whether that success is measured in terms of touchdown passes thrown, or in terms of patients surviving heart surgery, twenty percent of the people have fifty percent of the total successes.

Just think about it, what if we could make all our people, all our emergency managers, equivalent in their expertise over the next few years so they are at the top twenty percent. We would certainly find that we could do far more with fewer resources. It seems now that because of the budget deficit the next few years are likely to be times when we have to manage our resources more carefully than ever. This is the right time to look at expert systems. Expert systems will help us husband those resources.

Another legitimate question is: "Why should FEMA look at expert systems in emergency management?" Emergency management can be considered to be nearly an ideal field for the use of expert systems. The Department of Defense realizes this fact. They are leaders in this field, building a network of expert systems which they are calling their battlefield manager. They anticipate using expert systems in many ways to support their tactical and strategic systems in the future.

Fortunately most emergencies in the United States are on a much lower scale of intensity than those that the Department of Defense plans to face. But many of our problems are similar in nature. First of all, when an emergency occurs, it usually pays little attention to the expert's location or the time of day in which it occurs. Very often the person that is facing the emergency is not going to be an expert, nor will he always have direct access

to the needed expertise. Often we find that the person who is an expert is located in industry and is not immediately available to the government emergency manager. Expert systems may be able to fill these gaps in field operations.

In the area of emergency management we've taken the first steps towards identifying those functions for which we require expert systems. Since I began my tenure as the Director of the Federal Emergency Management Agency, a new policy has been adopted called the integrated emergency management system. Under this policy we no longer base all our planning on individual emergency situations. Rather we have identified certain functions which are common to any emergency.

These functions may be those of an emergency transportation manager, an emergency evacuation expert, or a warning expert. No matter what the emergency is, people have to be removed from the area of danger or protected in place. The injured have to be cared for. The temporarily homeless have to be housed and fed. The economic base of the affected region has to be rebuilt. These functions require high levels of expertise which are not always immediately or locally available. Therefore, the creation of expert systems in these various areas should help emergency managers to manage the limited resources available in a much more efficient and effective manner.

Another reason for investigating the utilization of expert systems in emergency management is that the environment is now right. As you look around you, you can find seven year olds using computers. Computer skills and the computers themselves are becoming more and more available. No longer do you have computer center gurus running large mainframe computers who talk to themselves and no one else. Rather you have many people that have the skills and abilities needed to interface with the computer. The technology itself is growing to the point that it is becoming natural to use computers.

Furthermore, as more and more people are using computers, the information flow available to the emergency manager is increasing. We in FEMA have been accommodating and encouraging this increase with the investment we've been making in computers and communications. Computer-based expert systems are a natural extension of this trend. And finally, no small factor in these days of tight budgets, Congress itself has expressed an interest in improving utilization of computers, communications, and other modern high technology techniques, in responding to emergencies.

So what is it that the Federal Emergency Management Agency expects from this symposium? First of all we want to listen to our emergency managers, to hear what they are saying, to hear their problems, to hear them tell us about those problems that they believe expert systems may be able to help. We also want to hear the experts on the development of expert systems. We want to hear what they say the solutions are, what they say they can do. We want to hear how they would respond to the problems that these emergency managers have advanced. From listening to these two groups we hope to develop a strategy for the future implementation of expert systems.

In essence, we want to develop a list of projects for the priority implementation of expert systems. We don't expect this conference to produce a panacea for emergency management. Nor do we expect a multitude of expert systems to spring up overnight. FEMA recognizes that this is a development project. However, we believe that the time to listen is now, the time to begin is today, the place to begin is here, listening to what you have to say. Thank you for your participation and your advice.

INTRODUCTORY REMARKS

Raymond G. Kammer
Deputy Director
National Bureau of Standards

It is with great pleasure that I welcome you on behalf of the National Bureau of Standards and our parent organization, the Department of Commerce, to this first symposium on the theory and application of expert systems in emergency management operations.

The challenge associated with the rapidly emerging field of expert systems is reminiscent of the beginnings of the information revolution. More than 30 years ago applied mathematicians at NBS used the SEAC -- Standards Eastern Automatic Computer -- to solve problems previously thought to be unsolvable. The first successful application of the now famous simplex algorithm was performed on the SEAC as a part of NBS's defense-related activities.

One of the first demonstrations of the value of computers is associated with research which led to the development of the hydrogen bomb. While a team of renowned physicists at Princeton were responsible for calculating the yield of the device, two NBS mathematicians, who were performing the back up calculations on the SEAC, produced results of startling accuracy. What is perhaps most remarkable is that this was done with a machine which had about the same capability as a sophisticated desk top personal computer.

As computers came of age and the mathematical models built to run on them became more complex, the role of measurement and standards for information exchange came to occupy a central position. The Federal Information Processing Standards, promulgated under the provisions of the Brooks Act, comprise key components of NBS's technical and coordinating role for improving the utilization of computers and their associated software within the Federal Government.

Both pure and applied research scientists were quick to realize that mathematical models are laboratories for the study of complicated scientific and engineering phenomena. Although mathematical models have always been a part of the scientific endeavor, the advent of computers changed forever the way in which models are used. Mathematical models have become important tools, not only for scientists and engineers, but for decision makers at all levels of government and industry.

The practical experience of our applied mathematicians and many of their colleagues with computer-based models points towards a need for guidelines in the planning, development, testing and appropriate use of such models. These practical concerns led to research on the subject of model evaluation. The research on model evaluation, for which NBS has been a focal point over the past decade, is complementary to the emerging field of expert systems. We at NBS, as all of you here today, are striving to explore the potential of these new techniques.

When we first began organizing this symposium, our goal was to establish a forum through which researchers in artificial intelligence could meet and interact with specialists in emergency management. As the symposium approached and our staff responded to the numerous inquiries, it became evident that many participants were interested not only in procedures for developing expert systems but for evaluating them as well.

NBS' work in this area should provide insight into the utility of expert systems in emergency management. These insights may reduce the challenge facing us, and they may increase the speed with which these new tools become a part of the decision-making process. In view of the urgency and complexity of the tasks confronting us, we anticipate that the work which will be described to you by our distinguished speakers over the next two days will be both stimulating and thought provoking.

A THEORY OF EXPERT SYSTEMS

H. Keith Hall, James C. Moore and Andrew B. Whinston

Krannert Graduate School of Management
Purdue University, West Lafayette, IN 47907

Introduction

The expert system approach to computer-based problem solving has received a great deal of attention in the last few years. Expert systems have been developed and implemented in such diverse areas as medical diagnosis, oil exploration, computer configuration, and logical data base design.¹ The most popular approach of current expert systems has been to represent expert knowledge within the system as a set of production rules that describes the process of decision making and characterizes the process that human experts go through in acquiring information and making decisions under uncertainty. Although there have been a number of different areas of research relating to expert systems including the theoretical basis for many of the computational techniques used, little research has been done on the decision theoretic basis for such problem solving systems in general. That is, they have not generally attempted to characterize an "ideal" process by which a decision could be made. Classical decision theory as used in economics, however, has typically characterized "best" decisions or decision strategies that may be formulated by a rational decision-maker. The basic goal of this paper, therefore, is to briefly discuss the decision-theoretic approach to decision-making under uncertainty and to compare this with that taken in many current expert systems. The major difference is the focus on the "correctness" of a decision and the efficiency of the decision process. An extension of the decision-theoretic approach to allow more complex information gathering strategies where the tradeoff between the cost of acquiring further information and the potential benefit from a better (more informed) decision made under less uncertainty may be found in Moore, Hall, and Whinston (1985).

The first section of the current paper presents the basic approach to acquiring information and making decisions under uncertainty taken in economics. The second section discusses current expert systems and presents a simple theoretical reformalization of this approach that is seen to be logically consistent in decision-making under uncertainty in economics. Both approaches employ information gathering strategies that generate an information structure upon which a decision strategy is based. In the last section, the expert knowledge necessary to construct an "ideal" expert system as would be suggested by the decision theoretic approach to decision-making under uncertainty is briefly discussed.

1. The Economics of Information

The economic approach to decision making has typically been concerned with the efficient choice of an alternative from a limited set of alternatives. The term efficiency has been broadly defined as the "best" decision made by a rational, utility maximizing decision-maker. When the decision is to be made under uncertainty, the decision-maker must make inferences about the relevant portions of his environment (the decision problem domain) after obtaining information

1. See Gevarter (1983) or Nau (1983) for briefs discussions of the techniques and many of the application areas of existing expert systems. Also see Bonczek, Holsapple, and Whinston (1981) for an exploration of management applications of expert systems.

about certain parameters that affect the outcome of his/her decision. These inferences are based upon a set of beliefs about the likelihood of the various possible states of the environment as embodied in a subjective probability distribution. The information that may be acquired prior to the decision, such as running experiments or test, asking questions, etc..., generally reduces uncertainty by altering this set of beliefs (likelihoods). The major difference in this approach from the current view of expert systems, however, is the focus on the "correctness" of a decision (finding the "best" solution) provided by the development in economics of the theory of preferences under uncertainty and the recognition that there is a tradeoff between correctness and the cost of acquiring the information necessary to improve the decision.

In general, economic models of decision-making under uncertainty by a single decision-maker may be placed into four of the following five categories. The fifth represents the extension suggested by the authors cited in the introduction above.

1) **Passive Response:**

- a) A decision is made without any information gathering, usually to maximize the expected payoff of the decision.
- b) A decision is to be made after receiving an (exogenous) information signal, so that a decision strategy (the decision made is contingent upon which signal is received) is formulated, usually to maximize the expected payoff of the decision.
- c) A number of decisions are to be made, each after receiving an information signal, so that a sequential decision strategy must be formulated, usually to maximize the expected utility of the decisions.

2) **Active Response:**

- a) A single information gathering action is chosen which generates an information signal. The action is chosen and a decision strategy is formulated usually to maximize the expected net payoff of the decision (taking the cost of the action into account).²
- b) A number, $r > 1$, of information gathering actions may be chosen, each generating an information signal. An information gathering strategy and a decision strategy are formulated. Each action and decision is contingent upon the signals received from previously executed actions in the sequence. Again, the objective is to maximize the expected net payoff of the decision.

1.1. The Basic Model of Decision Making

The environment is comprised of a finite set of states denoted by $X = \{x\}$. The decision-maker is assumed to have a set of feasible (final) decisions, D , and to receive a (gross) payoff which depends upon the state of the environment and the decision chosen. There is a deterministic relationship between decisions, states of the environment, and a set of outcomes (or results), O , such that there exists an outcome function, $o(x, d)$ mapping the set $X \times D$ into the outcomes. If the decision-maker's preferences over the outcomes and the cost of the decision may be represented by a real valued utility function, $u(o, d)$, for all $o \in O$ and $d \in D$, then the (gross) payoff function may be defined by

$$\omega(x, d) = u[o(x, d), d].$$

For the remainder of our theoretical discussion we will take the (gross) payoff function, $\omega(\cdot)$, as

2. This may also be interpreted as the choice among various information gathering *strategies*, each of which generates a sequence of information signals. Under this interpretation, this type of model would represent the "normal" form of the decision model in 2b) (the extensive form).

given.

While the payoff function alone may be used to order decisions according to the decision-maker's preferences given the knowledge of the true state of the environment, the connection between preferences and decisions under uncertainty is made by the "expected-utility" theorem of von Neuman and Morgenstern (1944). According to the theorem, a consistent decision-maker will not prefer decision $d' \in D$ to $d \in D$ if³

$$\sum_{x \in X} \omega(x, d) \Pr\{x\} \geq \sum_{x \in X} \omega(x, d') \Pr\{x\}.$$

Example. We will now use a simple example to illustrate the concept of expected utility (payoff). Let a decision-maker be choosing from among five different decisions and let there be eight different possible states of the environment. Further, there is the following gross payoff function

Payoff Function								
decisions	States							
	X ₁	X ₂	X ₃	X ₄	X ₅	X ₆	X ₇	X ₈
d ₁	32	28	24	20	28	24	16	12
d ₂	28	24	32	30	24	20	12	12
d ₃	24	16	20	16	32	24	32	30
d ₄	24	24	24	32	16	32	12	32
d ₅	16	32	12	28	8	28	0	20

and probability distribution

Pr{X ₁ }	Pr{X ₂ }	Pr{X ₃ }	Pr{X ₄ }	Pr{X ₅ }	Pr{X ₆ }	Pr{X ₇ }	Pr{X ₈ }
0	3/8	1/16	1/16	3/16	1/16	1/4	0

In comparing the five different decisions, an ordering may be formed with the equation

$\sum_{x \in X} \omega(x, d) \Pr\{x\}$ as follows:

d ₁	d ₃	d ₂	d ₄	d ₅
24	23 3/4	21 5/8	20 1/2	17 3/4

□

1.1.1. Information Acquisition

Associated with the states of the environment is a finite set of information gathering actions, which may be denoted without loss of generality by $A = \{0, 1, \dots, n\}$ and a cost function $c: A \rightarrow R$. Associated with each element of A is a set of $n(a)$ potential information messages which may, without loss of generality, be written $Y_a = \{1, 2, \dots, n(a)\}$. Each is seen to create a partition on X

3. For a brief discussion of the expected-utility theorem and decision making under uncertainty see Hirshleifer and Riley (1979) pp. 1379-1381.

and hence there implicitly exists the functions $\eta_a : X \rightarrow Y_a$ for $a \in A$ that may be used to define those states of the environment consistent with a given information signal $y \in Y_a$.

Definition 1.⁴ The function $\eta_a : X \rightarrow Y_a$ is called an information structure on X imposed by a .

We will let $a=0$ denote no additional information collection so that its associated partition is the entire state space.

Instead of the choice of a single information action, the decision maker may instead be choosing from among different information gathering strategies. That is, during the decision process a sequence of information gathering steps is taken, each contingent upon the results of previous steps, so that prior to the decision process a strategy must be formulated.⁵ To distinguish between the choice of single information gathering actions and the choice of different information strategies we will use $\alpha \in \Lambda$ to denote a strategy and continue to use $a \in A$ to denote a single action. For example, in figure 1 below, the decision tree representation of an information strategy, α_i , results in one of four different sequences of actions and seven different associated sequences of information signals. Each node represents a single information gathering action and each branch from a node represents a different possible signal. Also, each of the seven paths through the tree is associated with a sequence of signals.




Figure 1: Information Strategy α_i

Each step in the sequence of information actions is the continued partitioning of the state space into two subsets. One subset is known to contain the "true" state of the environment and the other is known not to. As more information is collected, the subset that contains the true state gets smaller and therefore uncertainty is reduced. For example, if a medical test is run by a physician, then the result is a signal that is known to be associated with certain diseases and not with others. Further, if a second medical test is run, then the second signal is consistent only with

4. An alternative definition of an information structure based on an underlying matrix of conditional probabilities relating the states to the signals and the marginal probabilities of the signals has been also used. See Marschak and Miyasawa (1968).

5. See 2a) and 2b) above in section 1.

portions of the current "possible" subset of the state space and therefore the addition of the new information reduces uncertainty as to the true disease. New updated likelihoods of the various states of the environment are formed using some known underlying probability distribution. Each strategy is therefore related to an information structure on X , $\eta_\alpha : X \rightarrow Y_\alpha$, where Y_α is the set of the possible sequences of information signals. This choice between different information strategies is the "normal" form of the decision problem (which may be thought of as a game against nature). In an "extensive" form of this same decision problem, found in Moore, Hall, and Whinston (1985), the focus is on the generation of an information structure rather than the choice of one. This would represent what classical decision theory might suggest is an "ideal" model of an expert system.⁶

1.1.2. Comparison of Information Structures

The value to a decision-maker of an information strategy and therefore its resulting information structure comes from the effect it's sequence of signals has on the choice of a final decision. The *ex post* value of an information strategy may therefore be measured by the change in the payoff due to the information minus the cost of the sequence of actions actually taken during the decision process

$$\omega(x, d_\alpha) - \omega(x, d_0) - \text{cost}$$

where d_α is the decision made after receiving the sequence $y \in Y_\alpha$ of information signals and d_0 is the decision that would have been made without additional information.

Since an information strategy must be chosen before a sequence of signals is received and a final decision is made, the decision-maker must rely instead upon the expected (*a priori*) value of an information strategy. Since a decision is contingent upon the sequence of information signals, a decision function must be first specified and we will denote this by

$$\delta : Y_\alpha \rightarrow D.$$

Following the expected-utility theorem the following definitions may then be made.

Definition 2. The maximum expected gross payoff of sequence y , the optimal decision rule for sequence y , and the maximum expected gross payoff of strategy α are defined by

$$v(y) = \max_{d \in D} \sum_{x \in X} \Pr(x | y) \omega(x, d)$$

$$= \sum_{x \in X} \Pr(x | y) \omega[x, \delta^*(y)]$$

$$\delta^*(y) = \operatorname{argmax}_{d \in D} \sum_{x \in X} \Pr(x | y) \omega(x, d)$$

$$\Omega(\alpha) = \sum_{y \in Y_\alpha} \Pr(y) v(y)$$

□

Since each strategy is associated with its resulting information structure, $\eta_\alpha(\cdot)$, $\Omega(\alpha)$ may be used to compare information structures. That is, η_α is said to be *more valuable* than $\eta_{\hat{\alpha}}$ if $\Omega(\alpha) \geq \Omega(\hat{\alpha})$.

Example cont. We now wish to illustrate the concept of the value of an information structure building upon our earlier example. Suppose the decision-maker may now gather information

6. See Owens (1983) pp. 1-6 for a general discussion of the normal and extensive forms of a game.

using one of five different information strategies and is therefore choosing from five different information structures and sequences of signal sets:

$$\eta_0(X) = 1 \text{ \& } Y_0 = \{1\};$$

$$\eta_1(x_1) = \eta_1(x_2) = 1, \eta_1(x_3) = \dots = \eta_1(x_8) = 2 \text{ \& } Y_1 = \{1, 2\};$$

$$\eta_2(x_1) = \eta_2(x_2) = 1, \eta_2(x_5) = \eta_2(x_6) = 2, \eta_2(x_3) = \eta_2(x_4) = \eta_2(x_7) = \eta_2(x_8) = 3 \\ \text{\& } Y_2 = \{1, 2, 3\};$$

$$\eta_3(x_1) = \eta_3(x_2) = 1, \eta_3(x_3) = \eta_3(x_4) = \eta_3(x_7) = \eta_3(x_8) = 2, \eta_3(x_5) = \eta_3(x_6) = 3 \\ \text{\& } Y_3 = \{1, 2, 3\};$$

$$\eta_4(x_1) = \eta_4(x_2) = 1, \eta_4(x_3) = \eta_4(x_4) = 2, \eta_4(x_5) = \dots = \eta_4(x_8) = 3 \text{ \& } Y_4 = \{1, 2, 3\}.$$

The information structure η_0 represents no information collection, η_1 partitions X into two sets, $\{x_1, x_2\}$ and $\{x_3, x_4, x_5, x_6, x_7, x_8\}$, η_2 partitions X into $\{x_1, x_2\}$, $\{x_5, x_6\}$, and $\{x_3, x_4, x_7, x_8\}$, η_3 partitions X into $\{x_1, x_2\}$, $\{x_3, x_4, x_7, x_8\}$, and $\{x_5, x_6\}$, and η_4 partitions X into $\{x_1, x_2\}$, $\{x_3, x_4\}$, and $\{x_5, x_6, x_7, x_8\}$.

Given the probability distribution for X , the marginal probability of each signal, $\Pr\{y\}$, and conditional probability of each state given each possible sequence of signals, $\Pr\{x | y\}$, may be calculated⁷ as well as an optimal decision rule for each sequence of signals. The optimal decision rules for each strategy are:

η_0	η_1	η_2	η_3	η_4
$\delta^*(1) = d_1$	$\delta^*(1) = d_5$ $\delta^*(2) = d_3$	$\delta^*(1) = d_5$ $\delta^*(2) = d_3$ $\delta^*(3) = d_3$	$\delta^*(1) = d_5$ $\delta^*(2) = d_3$ $\delta^*(3) = d_3$	$\delta^*(1) = d_5$ $\delta^*(2) = d_2$ $\delta^*(3) = d_3$

In comparing the five different information structures, an ordering may be formed as follows:

$\Omega(4)$	$\Omega(1)$	$\Omega(2)$	$\Omega(3)$	$\Omega(0)$
31 3/8	29 3/4	29 3/4	29 3/4	24

□

2. Rule-Based Expert Systems

Expert systems are designed to apply specialized or expert knowledge in performing difficult problem solving tasks. They differ from other computer-based systems in that they are designed to solve problems that normally require human expertise and intelligence. While there are rule-based, frame-based, and logic-based expert systems and most of our discussion may be applied to all three types of knowledge representation, we will focus on the first type that represents knowledge mainly as production rules.⁸ Rule-based expert systems are comprised of three basic components: a set of rules, a data base, and a rule interpreter. The data base temporarily holds the facts or assertions about a specific problem during the solution process. The rules of the rule set contain the domain-specific or expert knowledge. And the rule interpreter

7. See Marschak and Miyasawa (1968) pp. 140-142 for a discussion of the general properties of the standard probability measure and the calculation of conditional probabilities.

8. Frame-based systems tie knowledge objects and classes of objects connected by semantic links. Logic-based systems represent knowledge as assertions in logic, usually first order predicate logic and inference procedures are usually based on theorem proving.

contains the general problem-solving knowledge of the system that combines with the expert knowledge to solve problems. The interpreter functions by repeatedly matching production rules against data in the data base, selecting one or more of the rules with antecedent conditions met by the data (and therefore applicable to the particular problem being solved), and then executing the rule(s) which affects changes in the data. A solution is reached when an acceptable value is found for the variable or variables that represent the problem solution.




Figure 2: Components of a Rule-Based System

A production rule may be viewed as a conditional statement that relates factual statements with one another and are of the general form

IF <antecedent condition> THEN <action> WITH <CF>

The <antecedent condition> of a production rule is basically a conjunction of clauses that are either true or false. If the conjunction of clauses in the antecedent condition is currently true according to the database, then the <action> is executed which adds or replaces data in the database. When there is uncertainty in the validity of the rule, the <CF>, which stands for "certainty factor", represents the strength of belief in the rule. The knowledge embedded in the set of production rules are domain-specific. An example of a production rule from the MYCIN expert system is

ANTECEDENT CONDITION:

- (1) the infection is primary-bacteremia
- and
- (2) the site of the culture is one of the sterile sites
- and
- (3) the suspected portal of entry of the organism is the gastro-intestinal tract

ACTION:

there is suggestive evidence that the identity of the organism is bacteroides

WITH CF=0.7.

Since a solution is found by transforming the data in the database sequentially, the set of production rules may be formed into an AND/OR tree and the solution process represents the search of this tree for an acceptable solution. The path that the system takes through this tree represents the sequence of rules that are executed as determined by relevant data from the database. The rule interpreter controls this search for a solution so that it, with a given set of rules, defines a *solution strategy* that includes a strategy for acquiring information on each particular problem (*information strategy*) and for choosing a solution (*decision strategy*). The information and decision strategies may be sequential or nonsequential or both. The MYCIN medical diagnosis and treatment system, for example, goes through the following steps in interacting with the physician using the program:

- 1) The program begins with a set of questions about the patient. Since this same set of questions is asked regardless of the answers they may be asked in any order. This information acquisition is therefore nonsequential.
- 2) The program then asks for data on the patient as it is needed during the solution process to reduce the range of possible diseases. The information collected depends upon previously acquired information in step 1 and earlier in step 2 and is therefore collected in a specific order. The information acquired in step 2 is therefore collected sequentially. In applications to different patients, different sequences of questions will be asked in this step, depending upon the characteristics of the problem.
- 3) The program asks for the data on the patient needed to formulate a suggested treatment. This information is also sequential since it depends upon the information collected in step 2.
- 4) The program recommends treatment.

2.1. A Formal Representation

To discuss the solution strategy of an expert system we need to more formally present the basic expert system model. The data in the database may be represented by a set of parameters each with a number of possible values. In the example of the production rule from the MYCIN expert system presented in the previous subsection there are three parameters referred to in the antecedent condition and one in the action, each of which is assigned a value or range of values. If in the database the parameter representing infection is assigned the value denoting primary-bacteremia, the parameter representing the site of the culture is one of the sterile sites (i.e. the parameter falls within a range of values), and the parameter representing the portal of entry is assigned the value denoting the gastro-intestinal tract, then the value for the parameter presenting the identity of the organism is assigned the value that represents bacteroides. With the assumption that a production rule only refers to parameters in the database the parameters are of three types: those that may be given a value only by collecting information directly from a problem addressed by the system, those that may be assigned a value only by executing a production rule (therefore indirectly from a problem), and those that may be assigned a value either way (directly usually being the preferred way).

We will group these three types of parameters into two different sets. First, let there be a finite number K of parameters that represent those parameters that may be directly observed or measured from a problem. A state of the *observable environment* is therefore an assignment of a value to each of these parameters and we will represent this set of states by $Z = Z_1 \times Z_2 \times \cdots \times Z_K$ where Z_i is a set of possible values for the i^{th} observable parameter.

Second, let there be a finite number m of parameters that represent those parameters that may not be directly observed from a problem. A state of the *unobservable environment* is therefore an assignment of a value to each of these parameters and we will represent the set of states by $\Theta = \Theta_1 \times \Theta_2 \times \cdots \times \Theta_m$ where Θ_i is a set of possible values for the i^{th} unobservable parameter. An expert system, such as MYCIN, typically functions by sequentially reducing the range of possible observable states of the environment and then at some point in the sequence, selecting a final solution. During its solution process the MYCIN expert system uses a threshold of 0.2 (the likelihood that a particular state of the environment holds).⁹ Also, an expert system will, under certain values of the parameters in the database, collect a certain piece of information on a problem and under other values stop collecting information and select a certain problem solution. Since information collection involves only parameter values that may be directly observed, we may formulate two different sets of conditional statements or rules (not production rules) that are logically equivalent to the original set of production rules and the control strategy of the rule interpreter but involve only the observable states of the environment. First, there are *decision-making rules* which are conditional statements in which the conclusions assign a value to the parameter or parameters that represents a problem solution. Let there be a number K_1 of these rules and b_i represent a set of possible values for the data in the data base. We shall express these rules as:

IF b_i THEN d_i for $i=1, 2, \dots, K_1$.

The antecedent condition b_i of each conditional statement may be used to define a corresponding "truth set", which we will represent by B_i . This truth set is the set of all the possible states of the environment (i.e. $B_i \subseteq Z$) under which the conjunction of clauses in b_i (representing states of the database) are all true. Therefore, the antecedent condition b_i is seen to partition the states of the observable environment into two subsets: B_i and those states not in B_i .

It is important that an expert system correctly compute solutions for all problems in its problem domain. Knowledge base evaluation is the process of testing that a knowledge base is both complete and consistent. If a rule set is complete, then the rule processor may find a solution for all problems in the problem domain. If a rule set is consistent, then a consistent rule compiler will find the same solution each time a particular problem is solved by the system (or two different problems with the same relevant characteristics).¹⁰ Much of the research on expert systems has dealt with trying to insure consistency in rule-based expert systems where rules may be easily added or removed from a rule set and the rule interpreter must therefore often choose between several applicable rules during the solution process. For example, rule compilers are used frequently in expert systems and are designed to transform a rule set into decision tree form. The EMYCIN system, for example, has a rule compiler that transforms blocks of rules with common clauses into decision trees which evaluate several rules in parallel. For example, given the rules

A and B \rightarrow X

A and not B \rightarrow Y

not A and C \rightarrow Z

the decision tree in Figure 3 below would be formed by the rule compiler.¹¹

9. See pp. 225 in Davis, Buchanan, and Shortliffe (1977).

10. See Suwa, Scott, and Shortliffe (1984) for a general discussion of consistency and completeness in a rule set.

11. Taken from van Melle (1980) chapter 5.




Figure 3: A Compiled Block of Rules

Executing the decision tree representation is logically equivalent to interpreting the original block of rules but is more efficient since the redundancy in the antecedent conditions of the rules is removed and each clause is evaluated only once.

Using the above, we may more formally state the following two definitions of completeness and consistency.

Definition 3. The decision strategy of an expert system is complete if the union of the truth sets B_i , defined for each of the K_1 decision-making rules, equals the entire set of possible states of the environment. That is, the following holds:

$$Z = \bigcup_{i=1}^{K_1} B_i.$$

Definition 4. The decision strategy of an expert system is consistent if no two of the truth sets contain the same state of the environment. Or more formally, for any $i, j \in \{1, 2, \dots, K_1\}$, $i \neq j$ we have:

$$\text{IF } B_i \cap B_j \neq \emptyset \text{ THEN } d_i = d_j.$$

If one or both of the above conditions fail, then the rule set of the expert system is not correct. If a system is not complete, then there will be problems in the domain for which the system can reach no decision. Clearly, new decision-making rules need to be added until the system is complete. If a decision strategy is not consistent, then there will be problems in the domain for which the system reaches either more than one decision or different decisions on different applications to the same problem. The antecedent conditions of the existing rules need to be strengthened until the truth sets no longer overlap or a new rule needs to be added of the form:

$$\text{IF } b_i \text{ and } b_j \text{ THEN } d_{K_1+1}$$

and the overlapping rules must be modified to:

$$\text{IF } b_i \text{ and not } b_j \text{ THEN } d_i$$

If the decision strategy of an expert system is complete and consistent it is therefore seen to be equivalent to a decision function on the partition $\{B_1, B_2, \dots, B_{K_1}\}$ of the observable state space Z .

Second, there are *information-gathering rules* which are conditional statements in which the conclusions are in the form of a recommendation for acquiring further information on a specific problem. Let there be a number K_2 of these rules and we shall express them as:

IF b_i THEN a_i for $i=1, 2, \dots, K_2$.

where b_i represents a set of possible values for the data in the database and a_i is a request of further information on a problem. If the expert system were a medical diagnosis program, then this would be the recommendation that additional information on a patient be acquired by running an additional test(s) or asking an additional question(s).

If the information acquisition strategy of an expert system is consistent then the same sequence of actions will be executed for the same specific problem (with the same value of the observable critical parameters).¹² If the information strategy is complete then the sequence of actions always results in a decision. If an information strategy is both complete and consistent, then a decision tree may be formed with the actions and final decisions that represents the information and decision strategies of the system. To illustrate, in Figure 4 below the first tree represents both complete and consistent information and decision strategies, the second represents incomplete strategies, and the third represents inconsistent strategies. Each terminal node represents a decision (labeled with a d) and each of the other nodes represents an information action (labeled with an a). Each branch from a node represents a different possible result of the information action.




Figure 4: Solution Strategies

The set of information-gathering rules of an expert system is therefore seen to create the partition $\{B_1, B_2, \dots, B_{K_1}\}$ upon which the decision-making rules form their decisions. Since each $B_i \subseteq Z$ of a decision-making rule is associated with a sequence of information gathering steps (and therefore sequences of information signals) an information structure, $\eta_\alpha : Z \rightarrow Y_\alpha$, is generated by the information-gathering rules.

12. If the information strategy is nonsequential then this is trivially true since the same actions are executed for all problems.


2.2. Choice and Generation of Information Structures

Since a complete and consistent solution strategy generates an information structure different solution strategies represent possibly different choices of an information structures upon which to base a decision. The value of this choice is therefore measured by the correctness of the solutions generated by the system. In addition to the correctness of the solutions, the solution strategy also determines the efficiency of an expert system. A basic goal of expert system development is that this search for a solution be carried out selectively and intelligently. The concern with efficiency stems from the recognition that the solution process, particularly information collection, is costly and therefore costs should be taken into account. When a human expert goes through the process of collecting information prior to making a decision he/she implicitly makes judgements about the desirability of the possible outcomes and the cost of information. For example, when diagnosing and treating a patient, a physician will often decide that more tests should be run to find out if a patient has a relatively rare disease that, if incorrectly diagnosed and treated, is potentially very harmful. On the other hand, the same physician may decide not to collect further information on a patient even if he still has little idea which disease the patient suffers from, if he has eliminated all but those with the same best treatment or those that are not very harmful if misdiagnosed.

Example cont. Adding more detail to our earlier example, suppose the state of the environment consists of two different parts: the observable portion and the unobservable portion (i.e. $X = Z \times \Theta$). Let $Z = Z_1 \times Z_2 = \{0, 1\} \times \{0, 1\}$, $\Theta = \{\theta_1, \theta_2\}$,¹³ and the following partitions be created by each information action, $M_a = \{M_{a1}, M_{a2}\}$ for $a = 0, 1, \dots, 4$, where:

$$\begin{aligned} M_{11} &= \{z \in Z \mid z_1 + z_2 < 1\} & M_{12} &= \{z \in Z \mid z_1 + z_2 \geq 1\} \\ M_{21} &= \{z \in Z \mid z_2 = 0\} & M_{22} &= \{z \in Z \mid z_2 = 1\} \\ M_{31} &= \{z \in Z \mid z_1 - z_2 < 1\} & M_{32} &= \{z \in Z \mid z_1 - z_2 = 1\} \\ M_{41} &= \{z \in Z \mid z_1 = 0\} & M_{42} &= \{z \in Z \mid z_1 = 1\}. \end{aligned}$$

Now consider the following four strategies:



13. Note that $x_1 = (0, 0, \theta_1)$, $x_2 = (0, 0, \theta_2)$, ..., $x_7 = (1, 1, \theta_1)$, $x_8 = (1, 1, \theta_2)$.






Figure 5: Decision Tree Representations

and note that each generates one of the information structures, η_1, \dots, η_4 , described earlier. Let us further suppose that each information action has an associated cost, $c : A \rightarrow \mathbb{R}$, so that each

strategy now has an expected cost, $\gamma : \Lambda \rightarrow \mathbb{R}$. which may now be calculated as follows

$$\begin{aligned}\gamma(\alpha_1) &= c(1) \\ \gamma(\alpha_2) &= c(2) + \Pr(y=1)c(3) \\ &= c(2) + 3/8c(3) \\ \gamma(\alpha_3) &= c(3) + \Pr(y=1)c(1) \\ &= c(3) + 3/8c(1) \\ \gamma(\alpha_4) &= c(4) + \Pr(y=1)c(1) \\ &= c(4) + 3/8c(1)\end{aligned}$$

The maximum expected net payoff of strategy α is then defined as $\Omega(\alpha) - \gamma(\alpha)$ and given that $c(0)=0$, $c(1)=3$, $c(2)=4$, $c(3)=3$, and $c(4)=6$, may be used to create the following ordering on information structures:

$\Omega(1) - \gamma(1)$	$\Omega(3) - \gamma(3)$	$\Omega(2) - \gamma(2)$	$\Omega(4) - \gamma(4)$	$\Omega(0) - \gamma(0)$
26 3/4	25 5/8	24 5/8	24 1/4	24

□

3. The Acquisition of Expert Knowledge

There are three types of expert knowledge that must be acquired to implement an efficient expert system as described above. First is the knowledge of the problem domain that is necessary to identify a set of relevant parameters (and therefore the states of the environment) that describe each problem, identify a set of possible problem solutions (decisions), and calculate the *a priori* likelihoods of each of the various states of the environment. Second is the knowledge of the information gathering abilities of the system or system user which, combined with knowledge of the problem domain, is needed to update the likelihoods of the various possible states of the environment given the evidential support of one or more of the information actions. And lastly, knowledge of the desirability of the outcomes (and therefore the correctness) of the possible final decisions is needed to be weighed against the computational costs in guiding the solution process.

In the creation of an expert system, the acquisition of knowledge from an expert source generally begins with the conceptualization and structuring of knowledge of the domain. This is essentially the identification of many of the parameters and parameter values on which the production rules operate. Similarly, in classical decision theory under uncertainty a decision-maker is seen to make inferences about an uncertain environment after obtaining information about certain critical parameters that affect the consequences of his decisions. The importance of identifying these environmental parameters lies in their effect on the outcome of the decision and, more importantly, since the correctness of a decision is variable, in their effect on the desirability of the outcome. The identification of these parameters and their values cannot, therefore, be separated entirely from the identification of the different possible outcomes of a given solution and some sort of measure of their desirability.

The ease with which the environmental parameters may be identified in the construction of an expert system will vary with different problem domains. Generally, one would expect that expert knowledge that may easily be represented in productions rule form will also have a fairly developed set of critical parameters already identified by expert sources. An important distinction between a problem domain whose expert knowledge is easily represented by production rules and one whose expert knowledge is not, is whether it is best conceptualized as consisting of several independent states or by a single unified theory. Many examples of the former case may be found in medical diagnosis where expert sources have typically already identified and classified diseases, their associated symptoms, and possible treatments. The MYCIN expert

system, for example, diagnoses and suggests treatment for the class of infectious diseases. Problem domains that do not lend themselves well to representation by production rules are often best conceptualized not as various possible states but as a general unified theory that describes the relationships between possible states. Mathematics, for example, in representing knowledge about an operation such as addition, will focus on the property of commutativity that all operations called additivity share.¹⁴

If expert sources have already identified the parameters, then they will have done so with the set of possible decisions, the potential outcomes of the decisions, and the desirability of these outcomes at least implicitly in mind. In medical diagnosis, for example, it is obvious that distinguishing between two different diseases that typically either affect any given patient the same way no matter what the treatment (i.e. have the same effect on his/her well being) or have the same best treatment is less important to a physician than distinguishing between two other diseases with different effects on a patient or with different best treatments.

Expert knowledge of the domain must also be used to formulate *a priori* likelihoods (a frequency distribution) of the various states of the environment. An estimation of the likelihoods of the possible states is necessary because solutions are generally made without complete knowledge of the values of the critical parameters (i.e. without knowing the true state of the environment). The necessary data will normally be available in the form of frequency distributions over intervals and the intervals generally used in the presentation of the data in expert sources such as text books and papers will have an influence on the form of the specification of the environmental parameters. In the case of medical diagnosis, the incidence, within the general population to which the expert system is to be applied, of diseases of the various types which are to be diagnosed could be used. One useful expert source would be published public health data.

Knowledge of the system or system user's abilities to collect information on a specific decision problem is the second type of expert knowledge needed for an expert system. This, together with knowledge of the domain, allows the identification of a set of available information gathering actions and the calculation during the solution process of the likelihoods of the various states of the environment given the possible states of evidential support provided by previously executed information gathering actions. This knowledge is likely to be available from the same expert sources as knowledge of the domain. Therefore, if there is already a fairly well developed set of critical parameters for a particular domain, there will likely be a fairly developed set of information gathering actions. In medical diagnosis, for example, the standard medical tests used by physicians and data on the incidence within the population of the diseases in the domain associated with the different test results as well as estimates from human experts is generally available.

The last type of knowledge needed to implement our model as an expert system is knowledge of the desirability of the outcomes of the different solutions. The identification of the set of critical environmental parameters, set of possible solutions, and set of information actions already involves some of this knowledge because elements of these sets are important only through their effect on the outcome of the decision or more precisely on the desirability of the outcomes. The importance of using knowledge of the desirability of the possible outcomes to guide the solution process comes from the desire to efficiently reach problem solutions. In real applications, this is the recognition that collecting information is costly.

14. See pp. 29 Davis and King (1977)

References

- Bonczek, R.H., C.W. Holsapple and A.B. Whinston (1981). *Foundations of Decision Support Systems*, Academic Press, 1981.
- Buchanan, B.G. and R.O. Duda (1982). "Principles of Rule-Based Expert Systems", to appear in *Advances in Computers*, (M. Yovits, ed.), Vol. 22, August 1985.
- Davis, R., B.G. Buchanan, and E.H. Shortliffe (1977). "Production Rules as a Representation of a Knowledge-Based Consultant Program", *Artificial Intelligence*, 8, 15-45.
- Davis, R. and J.J. King (1984). "The Origin of Rule-Based Systems in AI", in *Rule-Based Expert Systems*, (B.G. Buchanan and E.H. Shortliffe, eds.), Addison-Wesley.
- Davis, R. et al., (1981). "The Dipmeter Advisor: Interpretation of Geological Signals", *Proc. Seventh Int'l Joint Conf. Artificial Intelligence*, August.
- Gevarter, W.B. (1983). "Expert Systems: Limited but Powerful", *IEEE*, August, 39-45.
- Gordon, J. and E.H. Shortliffe (1984). "The Dempster-Shafer Theory of Evidence", Chapter 13 in *Rule Based Expert Systems*, (B.B. Buchanan and E.H. Shortliffe, eds.), pp. 272-292.
- Hirshleifer, J. and R.G. Riley (1979). "The Analytics of Uncertainty and Information--An Expository Survey", *Journal of Economic Literature*, XVII, 1375-1421.
- Holsapple, C.W., S. Shen and A.B. Whinston (1982). "A Consulting System for Data Base Design", *International Journal Information Systems*, Vol. 7, no. 3, March, 281-296.
- Ishizuka, M., K.S. Fu, and J.T.P. Yao (1982). "Inference Procedures under Uncertainty for the Problem-Reduction Method", *Information Sciences*, 28, 179-206.
- Marschak, J. (1971). "Economics of Information Systems", in *Frontiers of Quantitative Economics* edited by M.D. Intriligator, Amsterdam: North-Holland, 32-107.
- Marschak, J. and K. Miyasawa (1968). "Economic Comparability of Information Systems", *International Economic Review*, 9, 137-174.
- McDermott, J. and B. Steele (1981). "Extending a Knowledge-Based System to Deal with Ad Hoc Constraints", *Proc. Seventh Int'l Joint Conf. Artificial Intelligence*, 824-828.
- van Melle, W. (1980). "A Domain-Independent System That Aids in Constructing Knowledge-Based Consultation Programs", Ph.D. Dissertation, Computer Science Department, Stanford University.
- Moore, J.C., H.K. Hall, and A.B. Whinston (1985). "Theoretical Foundations of Expert Systems", unpublished manuscript.
- Nau, D.S. (1983). "Expert Computer Systems", *IEEE*, February, 63-85.
- von Neumann, J. and O. Morgenstern (1944). *Theory of Games and Economic Behavior*, Princeton University Press, Princeton.
- Owen, G. (1983). *Game Theory* second edition, Academic Press, Inc.
- Shortliffe, E.H. and B.G. Buchanan (1975). "A Model of Inexact Reasoning in Medicine", *Mathematical Biosciences*, 23, 351-379.
- Shortliffe, E.H., R. Davis, S.G. Axline, B.G. Buchanan, C.C. Green, and S.N. Cohen (1975). "Computer-Based Consultations in Clinical Therapeutics: Explanation and Rule Acquisition Capabilities of the MYCIN System", *Computers and Biomedical Research*, 8, 303-320.
- Suwa, M., A.C. Scott, and E.H. Shortliffe (1984). "Completeness and Consistency in a Rule-Based System", Chapter 8 in *Rule Based Expert Systems*, (B.B. Buchanan and E.H. Shortliffe, eds.), pp. 159-170.

MANAGEMENT OF UNCERTAINTY IN EXPERT SYSTEMS

L. A. Zadeh^{*}

Abstract

During the past several years, the question of how to deal with uncertainty in the context of expert systems has attracted a great deal of attention because much of the information which is resident in the knowledge base of a typical expert system is imprecise, incomplete or not totally reliable.

The existing approaches to the management of uncertainty in expert systems are based for the most part on probability theory or its variants. However, it may be argued, as it is done in this paper, that probability theory is not sufficiently expressive as a language of uncertainty to represent the meaning of the imprecise facts and rules that form the knowledge base of a typical expert system. In an alternative approach which is outlined in this paper, fuzzy logic forms the basis for both meaning representation and inference. In particular, syllogistic reasoning is used to formulate a collection of rules for combination of evidence, with fuzzy quantifiers replacing probabilities and certainty factors as indicators of the degree of uncertainty.

1. Introduction

The issue of uncertainty plays an important role in the design and operation of expert systems because much of the information which is resident in the knowledge base of a typical expert system is imprecise, incomplete or not totally reliable.

The traditional probability-based methods of dealing with uncertainty do not work well in the case of expert systems because in practice it is usually infeasible to determine all of probability distributions which are needed to compute the conditional probability of a hypothesis given the evidence. This makes it necessary to employ ad hoc techniques of analysis in which many implicit and hard to verify assumptions regarding the underlying probability distributions are made.

Among the best known systems for dealing with uncertainty in the context of expert systems are those of MYCIN [11] and PROSPECTOR [18]. In MYCIN, for example, a typical rule is associated with a certainty factor, CF, which is a number in the interval $[-1,1]$. If positive, this number serves as a measure of the relative increase in the probability of a hypothesis given the evidence. If negative, its magnitude measures the relative decrease in the probability of the hypothesis given the evidence. The MYCIN rules of combination provide a system for computing the certainty factor of the conclusion from the certainty factors of the primary sources of evidence.

Although the MYCIN system works reasonably well in practice, it has a number of serious shortcomings which are analyzed in detail in a recent paper by Heckerman [23]. In his paper, Heckerman proposes several modifications

^{*} Computer Science Division, Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA 94720. Research supported in part by NASA Grant NCC2-275 and National Science Foundation Grant ECS-82 09679.

to the MYCIN system which make it much less ad hoc and are more solidly based on probability theory. However, there are some important sources of uncertainty in expert systems which are not addressed by the techniques based on probability theory and first-order logic. Among the more important of these sources are the following [52].

(1) The fuzziness of antecedents and/or consequents in rules of the form

(a) If X is A then Y is B,

(b) If X is A then Y is B with $CF = \alpha$,

where the antecedent, X is A, and the consequent, X is B, are fuzzy propositions, and α is a numerical value of the certainty factor, CF. For example [11]

(a) If the search space is moderately small then exhaustive search is feasible.

(b) If a piece of code is called frequently then it is worth optimizing.

(c) If large oil spill or strong acid spill then emergency is strongly suggested.

(d) If X is small then Y is large with $CF = 0.8$.

(e) If the route of the administration of the penicillin is oral, and there is a gastrointestinal factor which may interfere with the absorption of the penicillin, then there is suggestive evidence (0.6) that the route of administration of the penicillin is not adequate.

In these rules, the italicized words are labels of fuzzy predicates (e.g., small, large, emergency) or fuzzy predicate modifiers (e.g., moderately). This implies that the antecedents and consequents in the rules in question are fuzzy propositions or, equivalently, fuzzy events.

In the existing expert systems, the fuzziness of antecedents and consequents is ignored or treated incorrectly because neither probability-based methods nor bivalent logical systems provide a computational framework for dealing with it. As a consequence, fuzzy facts and rules are generally manipulated as if they were nonfuzzy, leading to conclusions of doubtful validity.

To illustrate this point, consider a rule of the general form [52]:

If X is A then Y is B with probability β ,

where X and Y are variables, A and B are fuzzy predicates and β is a fuzzy probability expressed as a fuzzy number, e.g., about 0.8, or as a linguistic probability, e.g., very likely. For example

If Mary is young then Mary is healthy is likely,

where $X \doteq \text{Age}(\text{Mary})$, $Y \doteq \text{Health}(\text{Mary})$, $A \doteq \text{young}$, $B \doteq \text{healthy}$ and $\beta \doteq \text{likely}$.

Expressed as a conditional probability, the rule in question may be written as

$$\Pr\{Y \text{ is } B | X \text{ is } A\} \text{ is } \beta. \quad (1.1)$$

In the existing expert systems, such a rule would be treated as an ordinary conditional probability, from which it would follow that

$$\Pr\{Y \text{ is not } B | X \text{ is } A\} \text{ is } 1 - \beta \quad (1.2)$$

However, as shown in [52], this conclusion is, in general, incorrect if A is a fuzzy set. The correct conclusion is weaker than (1.2), namely,

$$\Pr\{Y \text{ is not } B | X \text{ is } A\} + \Pr\{Y \text{ is } B | X \text{ is } A\} \geq 1, \quad (1.3)$$

with the understanding that the probabilities in question may be fuzzy numbers.

The above example points to two essential shortcomings of classical probability theory as a tool for dealing with uncertainty in expert systems. First, it makes no provision for fuzzy events; and second, it does not provide a mechanism for computing with fuzzy probabilities. Thus, the problem with the use of probability theory as a basis for the management of uncertainty in expert systems stems not from any flaws in its axiomatic foundations but from its lack of expressive power in representing the meaning of facts and rules which contain fuzzy predicates and/or fuzzy probabilities (or, equivalently, fuzzy quantifiers). Thus, viewed as a language, probability theory does not offer an appropriate framework for inference from propositions of the general form $X \text{ is } A$, if $X \text{ is } A$ then $Y \text{ is } B$, $X \text{ is } A \text{ is } \lambda$, if $X \text{ is } A$ then $Y \text{ is } B \text{ is } \lambda$, and QA 's are B 's, where A and B are fuzzy predicates (e.g., small, tall); λ is a fuzzy probability (e.g., likely, very unlikely), and Q is a fuzzy quantifier (e.g., most, many).

A view articulated in [52] is that the expressiveness of probability theory and predicate logic may be enhanced through the employment of fuzzy logic--a logical system which allows the use of fuzzy predicates and fuzzy quantifiers and thus subsumes both probability theory and predicate logic. In particular, through the use of fuzzy syllogisms, fuzzy logic provides an effective framework for the formulation of rules of construction of evidence in expert systems. In what follows, we shall present a summary of the basic concepts in fuzzy logic based on [50,52,53] and outline the rules of inference which are needed for dealing with uncertain facts and rules in the knowledge base of an expert system.

2. Fuzzy Logic

Like most logical systems, fuzzy logic has two components: (1) a representational component which deals with the representation of meaning of predicates, connectives, predicate modifiers and propositions; and (2) an inferential component, which, as its name implies, is concerned with the deduction of a conclusion from a set of premises. The representational component is based on PRUF [46] and test-score semantics [51], while the inferential component involves, for the most part, the manipulation of possibility distributions and fuzzy quantifiers.

More specifically, the first step in inference in fuzzy logic involves a translation of the premises p_1, \dots, p_n into PRUF--a meaning representation language which is based on the concept of a possibility distribution [46,49]. Thus, if p is a proposition and P is its translation into PRUF,

$$p \rightarrow P, \quad (2.1)$$

then P may be viewed as a procedure which acts on a collection of relations in a database--or, equivalently, a possible world--and returns a possibility distribution which represents the information conveyed by p.

If X is a variable taking values in U, then the possibility distribution of X, denoted by Π_X , is the fuzzy set of possible values of X. The membership function of Π_X is referred to as the possibility distribution function and is defined by

$$\pi_X(u) = \text{Poss}\{X = u\} \quad (2.2)$$

where the right-hand member of (2.2) should be read as: "The possibility that X can take u as a value," with the understanding that $0 \leq \pi_X(u) \leq 1$. More generally, if X is an n-ary variable

$$X = (X_1, \dots, X_n) \quad (2.3)$$

where X_i takes values in U_i , $i = 1, \dots, n$, then $\Pi_{(X_1, \dots, X_n)}$ is a fuzzy subset of the cartesian product $U \doteq U_1 \times \dots \times U_n$, and the projection of $\Pi_{(X_1, \dots, X_n)}$ on $U_{(s)} \doteq U_{i_1} \times \dots \times U_{i_k}$ is defined by the possibility distribution function

$$\begin{aligned} & X_{i_1} \times \dots \times X_{i_k} \pi_{(X_1, \dots, X_n)}(u_{i_1}, \dots, u_{i_k}) \\ &= \sup_{u_{j_1}, \dots, u_{j_\ell}} \pi_{(X_1, \dots, X_n)}(u_1, \dots, u_n) \end{aligned} \quad (2.4)$$

where $X_{(s)} \doteq (X_{i_1}, \dots, X_{i_k})$ is a subvariable of $X \doteq (X_1, \dots, X_n)$ and $X_{(s')} \doteq (X_{j_1}, \dots, X_{j_\ell})$ is the complementary subvariable, i.e., $\{j_1, \dots, j_\ell\} = \{1, \dots, n\} - \{i_1, \dots, i_k\}$.

An n-ary possibility distribution is particularized by forming the conjunction of the propositions "X is F" and " $X_{(s)}$ is G," where $X_{(s)}$ is a subvariable of X and F and G are fuzzy subsets of U. Thus,

$$\Pi_X[\Pi_{X_{(s)}} = G] \doteq F \cap \bar{G} \quad (2.5)$$

where the right-hand member denotes the intersection of F with the cylindrical extension of G, i.e., a cylindrical fuzzy set defined by

$$\mu_{\bar{G}}(u_1, \dots, u_n) = \mu_G(u_{i_1}, \dots, u_{i_k}), \quad (u_1, \dots, u_n) \in U_1 \times \dots \times U_n \quad (2.6)$$

where μ_G denotes the membership function of G.

There are many cases in which the operations of particularization and projection are combined. In such cases it is convenient to use the simplified notation

$$X_{i_1} \times \dots \times X_{i_k} \Pi[\Pi_{X_{(s)}} = G] \quad (2.7)$$

to indicate that the particularized possibility distribution (or relation) $\Pi[\Pi_{X(s)} = G]$ is projected on $U_{i_1} \times \dots \times U_{i_k}$. For example,

$$X_1 \times X_3 \Pi[\Pi_{(X_3, X_4)} = G] \quad (2.8)$$

would represent the projection of $\Pi[\Pi_{(X_3, X_4)} = G]$ on $U_1 \times U_3$. Informally, (2.7) may be interpreted as: Constrain the $X_{(s)}$ by $\Pi_{X(s)} = G$ and read out the X_i . In particular, if the value of $X_{(s)}$ --rather than its possibility distribution--is set equal to G, then (2.7) becomes

$$X_{i_1} \times \dots \times X_{i_k} \Pi[X_{(s)} = G]. \quad (2.9)$$

Now let p be a proposition of the form

$$p \doteq N \text{ is } F \quad (2.10)$$

where N is the name of an object, a variable or a proposition, and F is a fuzzy subset of a universe of discourse U. For example:

$$\begin{aligned} p &\doteq \text{Sheila is very intelligent} \\ q &\doteq X \text{ is small} \\ r &\doteq (X \text{ is small}) \text{ is more or less true} \end{aligned}$$

A translation of p may be focused or unfocused. By a focused translation in PRUF is meant a translation of the form

$$p \rightarrow \Pi_{(X_1, \dots, X_n)} = F \quad (2.11)$$

where the X_i are variables which are explicit or implicit in N and F is a fuzzy subset of $U_1 \times \dots \times U_n$. A translation of this form serves to (a) identify the variables in N whose possibility distribution is determined by p, and (b) specify the possibility distribution in question. To reflect this aspect of the translation process, the right-hand member of (2.11) is referred to as the possibility assignment equation, with F representing the possibility distribution induced by p.

As a simple illustration of (2.11), consider the proposition

$$p \doteq \text{Mary has blue eyes.} \quad (2.12)$$

In this case, (2.11) becomes

$$\text{Mary has blue eyes} \rightarrow \Pi_{\text{Color}(\text{Eyes}(\text{Mary}))} = \text{BLUE} \quad (2.13)$$

in which $X \doteq \text{Color}(\text{Eyes}(\text{Mary}))$ and BLUE, the denotation of blue, is the fuzzy subset of colors which are perceived as blue.

Turning to the concept of an unfocused translation, let $\mathcal{D} = \{D\}$ denote a collection of databases, with D representing a generic element of \mathcal{D} . For

the purposes of our analysis, D will be assumed to consist of a collection of possibly time-varying relations. If R is a constituent relation in D , then by the frame of R is meant the name of R together with the names of its columns (i.e., attributes). For example, if a constituent of D is a relation labeled POPULATION whose tableau is comprised of columns labeled Name and Height, then the frame of POPULATION is represented as POPULATION[Name,Height].

In relation to D , the unfocused translation of p can assume one of two equivalent forms:¹

- (a) $p \rightarrow$ a procedure which yields for each D in D the possibility of D given p , i.e., $\text{Poss}\{D|p\}$
- (b) $p \rightarrow$ a procedure which yields for each D in D the truth-value of p relative to D , i.e., $\text{Tr}\{p|D\}$

The equivalence of (a) and (b) is a consequence of the way in which the concept of truth is defined in fuzzy logic [53]. Thus, it can readily be shown that, under mildly restrictive assumptions on D , we have

$$\text{Tr}\{p|D\} = \text{Poss}\{D|p\} \quad (2.14)$$

which implies the equivalence of (a) and (b). Furthermore, it should be noted that (2.11) is a special case of (a), so that the concept of a focused translation is subsumed by that of an unfocused translation. In essence, our motivation for introducing the concept of a focused translation is to provide a mechanism for enhancing the explanatory effectiveness of a translation. In the case of a focused translation, this is accomplished by focusing on those and only those variables in D whose possibility distribution is constrained by the meaning of p .

If Π^p and Π^q are the possibility distributions induced by p and q , respectively, then p and q are semantically equivalent if and only if $\Pi^p = \Pi^q$.² In symbols

$$p \leftrightarrow q \text{ iff } \Pi^p = \Pi^q \quad (2.15)$$

For example, it can readily be shown that

$$X \text{ is small is true} \leftrightarrow X \text{ is not small is false}$$

provided the linguistic truth-value false is taken to be the antonym of true, i.e.,

$$\mu_{\text{false}}(v) = \mu_{\text{true}}(1-v), \quad 0 \leq v \leq 1 \quad (2.16)$$

where μ_{false} and μ_{true} are the membership functions of false and true,

¹It should be noted that (a) and (b) are in the spirit of possible-world semantics and truth-conditional semantics, respectively. In their conventional form, however, these semantics have no provision for fuzzy propositions and hence are not suitable for use in fuzzy logic.

²To be more precise, we have to differentiate between the concepts of semantic equivalence and strong semantic equivalence. The latter concept reduces to that of semantic equivalence in predicate logic when p and q are nonfuzzy propositions.

respectively, More generally, q is semantically entailed by p if and only if $\Pi^p \subset \Pi^q$. In symbols

$$p \vdash q \text{ iff } \Pi^p \subset \Pi^q. \quad (2.17)$$

For example,

$$X \text{ is very small} \vdash X \text{ is small} \quad (2.18)$$

if very is defined by

$$\mu_{\text{very small}}(u) = (\mu_{\text{small}}(u))^2, \quad u \in U. \quad (2.19)$$

We shall make use of the concept of semantic entailment in Section 3.

An essential component of a fuzzy logic is the set of translation rules, which may be categorized into four basic types: Type I - rules pertaining to modification; Type II - rules pertaining to composition; Type III - rules pertaining to quantification; and Type IV - rules pertaining to qualification. For our purposes, it will suffice to state the rules of Type II, which are:

$$\text{If } X \text{ is } F \rightarrow \Pi_X = G \text{ and } Y \text{ is } G \rightarrow \Pi_Y = G \quad (2.20)$$

where F and G are fuzzy subsets of U and V , respectively, then

$$(a) \text{ } X \text{ is } F \text{ and } Y \text{ is } G \rightarrow \Pi_{(X,Y)} = F \times G \quad (2.21)$$

$$\text{where } \mu_{F \times G}(u,v) = \mu_F(u) \wedge \mu_G(v). \quad (2.22)$$

$$(b) \text{ } X \text{ is } F \text{ or } Y \text{ is } G \rightarrow \Pi_{(X,Y)} = \bar{F} \cup \bar{G} \quad (2.23)$$

$$\text{where } \bar{F} \triangleq F \times V, \quad \bar{G} \triangleq U \times G \quad (2.24)$$

$$\text{and } \mu_{\bar{F} \cup \bar{G}}(u,v) = \mu_F(u) \vee \mu_G(v). \quad (2.25)$$

$$(c) \text{ If } X \text{ is } F \text{ then } Y \text{ is } G \rightarrow \Pi_{(Y|X)} = \bar{F}' \oplus \bar{G} \quad (2.26)$$

where $\Pi_{(Y|X)}$ denotes the conditional possibility distribution of Y given X , and the bounded sum \oplus is defined by

$$\mu_{\bar{F}' \oplus \bar{G}}(u,v) = 1 \wedge (1 - \mu_G(u) + \mu_G(v)). \quad (2.27)$$

It should be noted that the compositional rules in FL are defined in a way that makes them consistent with the valuation rules in Lukasiewicz's L_{Aleph1} logic. In particular, if $X = Y$ then it follows at once from (2.27) that the proposition

$$r \triangleq \text{If } X \text{ is } F \text{ then } X \text{ is } G \quad (2.29)$$

conveys no information concerning X (i.e., is a tautology) if and only if $F \subset G$. This conclusion serves as a basis for the entailment principle which is stated in the following section.

3. Rules of Inference

The rules of inference in fuzzy logic are semantic rather than syntactic in nature and are based, in the main, on two principles: (a) the entailment principle, and (b) the extension principle.

The entailment principle, which is based in turn on (2.28), may be expressed in a schematic form as

$$\begin{array}{c} p \rightarrow \Pi_X^p = F \\ \hline q \leftarrow \Pi_X^q = G: F \subset G \end{array} \quad (3.1)$$

which means that if the possibility distribution induced by p is contained in that induced by q , then q may be inferred from p . (The symbol \leftarrow stands for "retranslation.") In other words, from p we can infer any proposition which is semantically equivalent to p or, more generally, semantically entailed by p .

To state the extension principle, let f be a function from U or V and let F be a finite fuzzy subset of U expressed as

$$F = \mu_1/u_1 + \dots + \mu_n/u_n. \quad (3.2)$$

Then, the extension of f to the set of fuzzy subsets of U is defined by³

$$\begin{aligned} f(F) &= f(\mu_1/u_1 + \dots + \mu_n/u_n) \\ &= \mu_1/f(u_1) + \dots + \mu_n/f(u_n). \end{aligned} \quad (3.3)$$

More generally, if the support of F is a continuum, i.e.,

$$F = \int_U \mu_F(u)/u \quad (3.4)$$

then

$$f(F) = \int_U \mu_F(u)/f(u). \quad (3.5)$$

Furthermore, if U is a cartesian product of U_1, \dots, U_n and f is a mapping from $U_1 \times \dots \times U_n$ to V , then

$$f(F) = \int_U \mu_F(u_1, \dots, u_n)/f(u_1, \dots, u_n). \quad (3.6)$$

In particular, if we have only partial information about μ_F , e.g., the knowledge of its projections on U_1, \dots, U_n , then

$$f(F) = \int_U \mu_1(u_1) \wedge \dots \wedge \mu_n(u_n)/f(u_1, \dots, u_n) \quad (3.7)$$

³The notation $F = \mu_1/u_1 + \dots + \mu_n/u_n$ signifies that F is a collection of fuzzy singletons μ_i/u_i , $i = 1, \dots, n$, with μ_i representing the grade of membership of u_i in F . More generally, F may be expressed as $F = \sum_i \mu_i/u_i$ or

$$F = \int_U \mu_F(u)/u.$$

where μ_i , $i = 1, \dots, n$, is the membership function of the projection of F on U_i .

A more explicit representation of $f(F)$ which is equivalent to (3.7) is the following. ($\mu \triangleq$ membership function of $f(F)$.)

$$f(F) = \int_V \mu(v)/v \quad (3.8)$$

$$\text{where} \quad \mu(v) \triangleq \max_{u_1, \dots, u_n} \mu_1(u_1) \wedge \dots \wedge \mu_n(u_n) \quad (3.9)$$

subject to the constraint

$$v = f(u_1, \dots, u_n) .$$

In this form, the extension principle serves to reduce an inference process to the determination of the solution of a variational problem in nonlinear programming.

Among the rules of inference which are based on the entailment principle are the following.

Projection rule. Stated in a schematic form, this rule may be expressed as

$$\frac{p \rightarrow \Pi^p(x_1, \dots, x_n) = F}{q \leftarrow \Pi^q(x_{i_1}, \dots, x_{i_k}) = x_{i_1} \times \dots \times x_{i_k}^F} \quad (3.10)$$

where $x_{i_1} \times \dots \times x_{i_k}^F$ denotes the projection of F on $U_{i_1} \times \dots \times U_{i_k}$ (see (2.4)).

In other words, (3.10) asserts that if p induces the possibility distribution $\Pi^p(x_1, \dots, x_n)$, then from p we can infer any proposition q which induces the projection of Π^p on $U_{i_1} \times \dots \times U_{i_k}$.

As a simple illustration of (3.10), consider the proposition

$$p \triangleq \text{John is big}$$

and assume that BIG, the denotation of big, is a binary fuzzy relation which may be represented as the cartesian product of the unary fuzzy relations TALL and FAT. Then

$$\text{TALL} = \text{Height}^{\text{BIG}}$$

and hence we can assert that

$$\frac{\text{John is big}}{\text{John is tall}}$$

Conjunction rule. Let X , Y and Z take values in U , V and W , respectively; let F and G be fuzzy subsets of $U \times V$ and $V \times W$; and let \bar{F} and \bar{G} be the cylindrical extensions of F and G in $U \times V \times W$. Thus,

$$\bar{F} = F \times W$$

and

$$\bar{G} = U \times G.$$

The conjunction rule is defined by the following scheme:

$$\begin{array}{l} p \rightarrow \Pi_{(X,Y)}^p = F \\ q \rightarrow \Pi_{(Y,Z)}^q = G \\ \hline r \leftarrow \Pi_{(X,Y,Z)} = \bar{F} \cap \bar{G} \end{array} \quad (3.11)$$

In particular, in the special case where

$$\begin{array}{l} p \rightarrow \Pi_X^p = F \\ q \rightarrow \Pi_X^q = G \end{array} \quad (3.12)$$

(3.11) implies that from p and q we can infer r, where

$$\Pi_X^r = F \cap G.$$

Compositional rule. On combining the projection and conjunction rules, we obtain the compositional rule of inference (3.13) which includes the classical modus ponens as a special case.

More specifically, on applying the projection rule to (3.11), we obtain the following inference scheme

$$\begin{array}{l} p \rightarrow \Pi_{(X,Y)}^p = F \\ q \rightarrow \Pi_{(Y,Z)}^q = G \\ \hline r \leftarrow \Pi_{(X,Z)}^r = F \circ G \end{array} \quad (3.13)$$

where the composition of F and G is defined by

$$\mu_{F \circ G}(u,w) = \sup_v (\mu_F(u,v) \wedge \mu_G(v,w)). \quad (3.14)$$

In particular, if p is a proposition of the form "X is F" and q is a proposition of the form "If X is G then Y is H," then (3.13) becomes

$$\begin{array}{l} p \rightarrow \Pi_X = F \\ q \rightarrow \Pi_{(Y|X)} = \bar{G}' \oplus \bar{H} \\ \hline r \leftarrow \Pi_Y = F \circ (\bar{G}' \oplus \bar{H}) \end{array} \quad (3.15)$$

The rule expressed by (3.15) may be viewed as a generalized form of modus ponens which reduces to the classical modus ponens when $F = G$ and F, \bar{G}, \bar{H} are nonfuzzy sets.

Semantic equivalence. An immediate consequence of the entailment principle is that from p we can infer any proposition q which is semantically equivalent to p .

An important special case of this rule is embodied in the truth-qualification principle which may be stated as follows.

Let τ be a linguistic truth-value which is characterized by its membership function μ_τ and assume that

$$p \doteq X \text{ is } F.$$

Then

$$X \text{ is } F \leftrightarrow X \text{ is } G \text{ is } \tau \quad (3.16)$$

where the membership functions of F , G and τ are related by

$$\mu_F(u) = \mu_\tau(\mu_G(u)). \quad (3.17)$$

Thus, from the proposition $q \doteq X \text{ is } G \text{ is } \tau$ we can infer $p \doteq X \text{ is } F$, and vice-versa, with the understanding that μ_F is given by the composition of μ_τ and μ_G .

As an illustration, consider the proposition

$$p \doteq \text{Lynn is young is very true}$$

in which

$$\mu_{\text{young}}(u) = (1 + (\frac{u}{20})^2)^{-1}, \quad u \in [0, 100]$$

and

$$\mu_{\text{true}}(v) = v^2, \quad v \in [0, 1].$$

Then, we can infer from p that

$$\pi_{\text{Age(Lynn)}}(u) = (1 + (\frac{u}{20})^2)^{-4}$$

which upon retranslation may be expressed as

$$q \doteq \text{Lynn is very very young.}$$

Transformational rule. This rule is implied by the extension principle and may be expressed in a schematic form as

$$\begin{array}{l} p \doteq X \text{ is } F \\ q \doteq Y \text{ is } f(X) \\ \hline r \doteq Y \text{ is } f(F) \end{array} \quad (3.18)$$

where F is a fuzzy subset of U , f is a mapping from U to V , and $f(F)$ is a fuzzy subset of V defined by (3.8).

As a simple illustration of (3.18), assume that X is a fuzzy number which is expressed as "approximately 5" or, more explicitly,

$$X = 0.2/2 + 0.5/3 + 0.8/4 + 1/5 + 0.8/6 + 0.6/7 + 0.2/8$$

and $Y = X^2$. Then

$$Y = 0.2/4 + 0.5/9 + 0.8/16 + 1/25 + 0.8/36 + 0.6/49 + 0.2/64.$$

4. Syllogistic Reasoning

In addition to the rules described in the preceding section, there is a class of rules of inference in fuzzy logic which are of particular relevance to the combination of evidence in expert systems. The rules in question may be viewed as fuzzy syllogisms [54] in which the premises are propositions of the form QA's are B's, where A and B are fuzzy predicates and Q is a fuzzy quantifier.

In its generic form, a fuzzy syllogism may be expressed as the inference schema

$$\begin{array}{l} Q_1 A's \text{ are } B's \\ Q_2 C's \text{ are } D's \\ \hline Q_3 E's \text{ are } F's \end{array} \quad (4.1)$$

in which A, B, C, D, E and F are interrelated fuzzy predicates and Q_1 , Q_2 and Q_3 are fuzzy quantifiers.

The interrelations between A, B, C, D, E and F provide a basis for a classification of fuzzy syllogisms. The more important of these syllogisms are the following. ($\wedge \doteq$ conjunction, $\vee \doteq$ disjunction).

- (a) Intersection/product syllogism: $C = A \wedge B$, $E = A$, $F = C \wedge D$
- (b) Chaining syllogism: $C = B$, $E = A$, $F = D$
- (c) Consequent conjunction syllogism: $A = C = E$, $F = B \wedge D$
- (d) Consequent disjunction syllogism: $A = C = E$, $F = B \vee D$
- (e) Antecedent conjunction syllogism: $B = D = F$, $E = A \wedge C$
- (f) Antecedent disjunction syllogism: $B = D = F$, $E = A \vee C$.

In the context of expert systems, these and related syllogisms provide a set of inference rules for combining evidence through conjunction, disjunction and chaining [52].

One of the basic problems in fuzzy syllogistic reasoning is the following. Given A, B, C, D, E and F, find the maximally specific (i.e., most restrictive) fuzzy quantifier Q_3 such that the proposition $Q_3 E's \text{ are } F's$ is entailed by the premises. In the case of (a), (b) and (c), this leads to the following syllogisms.

Intersection/product syllogism

$$\begin{array}{l} Q_1 A's \text{ are } B's \\ Q_2 (A \text{ and } B)'s \text{ are } C's \\ \hline (Q_1 \otimes Q_2) A's \text{ are } (B \text{ and } C)'s \end{array} \quad (4.2)$$

where \otimes denotes the product in fuzzy arithmetic [25]. It should be noted that (4.2) may be viewed as an analog of the basic probabilistic identity $p(B, C/A) = p(B/A)p(C/A, B)$.

A concrete example of the intersection/product syllogism is the following:

$$\begin{array}{l} \text{most students are young} \\ \text{most young students are single} \\ \hline \text{most}^2 \text{ students are young and single,} \end{array} \quad (4.3)$$

where most^2 denotes the product of the fuzzy quantifier most with itself.

Chaining syllogism

$$\begin{array}{l} Q_1 \text{ A's are B's} \\ Q_2 \text{ B's are C's} \\ \hline (Q_1 \otimes Q_2) \text{ A's are C's.} \end{array} \quad (4.4)$$

This syllogism may be viewed as a special case of the intersection/product syllogism. It results when $B \subset A$ and Q_1 and Q_2 are monotone increasing, i.e., $\geq Q_1 = Q_1$, and $\leq Q_2 = Q_2$, where $\geq Q_1$ should be read as at least Q_1 , and likewise for Q_2 . A simple example of the chaining syllogism is the following:

$$\begin{array}{l} \text{most students are undergraduates} \\ \text{most undergraduates are single} \\ \hline \text{most}^2 \text{ students are single} \end{array}$$

Note that undergraduates \subset students and that in the conclusion $F = \text{single}$, rather than young and single, as in (4.3).

Consequent conjunction syllogism. The consequent conjunction syllogism is an example of a basic syllogism which is not a derivative of the intersection/product syllogism. Its statement may be expressed as follows:

$$\begin{array}{l} Q_1 \text{ A's are B's} \\ Q_2 \text{ A's are C's} \\ \hline Q \text{ A's are (B and C)'s,} \end{array} \quad (4.5)$$

where Q is a fuzzy quantifier which is defined by the inequalities

$$0 \otimes (Q_1 \oplus Q_2 \ominus 1) \leq Q \leq Q_1 \otimes Q_2 \quad (4.6)$$

in which \otimes , \oplus , \ominus and \ominus are the operations of \vee (max), \wedge (min), $+$ and $-$ in fuzzy arithmetic.

An illustration of (4.5) is provided by the example

most students are young
most students are single
 Q students are single and young

where $2\text{most} \ominus 1 \leq Q \leq \text{most}.$ (4.7)

This expression for Q follows from (4.6) by noting that

$$\text{most} \otimes \text{most} = \text{most}$$

and

$$0 \otimes (2\text{most} \ominus 1) = 2\text{most} \ominus 1.$$

The three basic syllogisms described above are representative of a class of rules of inference in fuzzy logic in which fuzzy quantifiers play a role analogous to that of probabilities in probabilistic reasoning. Through the use of such rules, probabilistic reasoning becomes a part of fuzzy logic. In this way, it becomes possible to employ a single conceptional framework for the management of uncertainty in expert systems in which both probability theory and possibility theory play central roles.

References and Related Publications

- [1] K. P. Adlassnig, A survey of medical diagnosis and fuzzy subsets, in: M. M. Gupta and E. Sanchez, Eds. Approximate Reasoning in Decision Analysis (North-Holland, Amsterdam, 1982) 203-217.
- [2] J. F. Baldwin, A new approach to approximate reasoning using a fuzzy logic, Fuzzy Sets and Systems 2 (1979) 302-325.
- [3] W. Bandler, Representation and manipulation of knowledge in fuzzy expert systems, Proc. Workshop on Fuzzy Sets and Knowledge-Based Systems, Queen Mary College, University of London (1983).
- [4] A. Barr and E. W. Feigenbaum, The Handbook of Artificial Intelligence, Vol. 1, 2 and 3 (Kaufmann, Los Altos, 1982).
- [5] J. Barwise and R. Cooper, Generalized quantifiers and natural language, Linguistics and Philosophy 4 (1981) 159-219.
- [6] R. E. Bellman and L. A. Zadeh, Local and fuzzy logics, in: G. Epstein, Ed. Modern Uses of Multiple-Valued Logic (Reidel, Dordrecht, 1977) 103-165.
- [7] M. Ben-Bassat, et al., Pattern-based interactive diagnosis of multiple disorders: The MEDAS system, IEEE Trans. Pattern Anal. Machine Intell. (1980) 148-160.
- [8] J. Bezdek, Pattern Recognition with Fuzzy Objective Function Algorithms (Plenum Press, New York, 1981).

- [9] P. P. Bonissone and R. M. Tong, Eds., Special Issue on Reasoning with Uncertainty in Expert Systems, *Int. J. Man-Machine Studies* 22 (1985).
- [10] R. J. Brachman and B. C. Smith, Special Issue on Knowledge Representation, *SIG ART* 70 (1980).
- [11] B. B. Buchanan and E. H. Shortliffe, *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project* (Addison-Wesley, Reading, MA, 1984).
- [12] P. Cheeseman and J. Lemmer, Eds., *Proceedings of Workshop on Certainty and Probability in Artificial Intelligence*, UCLA, Los Angeles (1985).
- [13] M. S. Cohen, D. A. Schum, A. N. S. Freeling and J. O. Chinnis, On the art and science of hedging a conclusion: alternative theories of uncertainty in intelligence analysis, *Tech. Rep. 84-6*, Decision Science Consortium, Falls Church, VA, 1984.
- [14] R. Davis and D. B. Lenat, *Knowledge-Based Systems in Artificial Intelligence* (McGraw-Hill, New York, 1982).
- [15] A. DeLuca and S. Termini, A definition of non-probabilistic entropy in the setting of fuzzy sets theory, *Inform. and Control* 20 (1972) 301-312.
- [16] D. Dubois and H. Prade, *Fuzzy Sets and Systems: Theory and Applications* (Academic Press, New York, 1980).
- [17] D. Dubois and H. Prade, On several representations of an uncertain body of evidence, in M. M. Gupta and E. Sanchez, Eds., *Fuzzy Information and Decision Processes* (North-Holland, Amsterdam, 1982) 167-181.
- [18] R. O. Duda, P. E. Hart, K. Konolige and R. Reboh, A computer-based consultant for mineral exploration, *Final Tech. Report*, SRI International, Menlo Park, CA (1979).
- [19] E. A. Feigenbaum, The art of artificial intelligence: Themes and case studies in knowledge engineering, *Proc. IJCAI5* (1977).
- [20] B. R. Gaines, Logical foundations for database systems, *Int. J. Man-Machine Studies* 11 (1979) 481-500.
- [21] T. Garvey and J. Lowrance, Evidential reasoning: implementation for multisensor integration, *SRI International Tech. Rep. 307*, Menlo Park, CA (1983).
- [22] J. A. Goguen, The logic of inexact concepts, *Synthese* 19 (1959) 325-373.
- [23] D. Heckerman, Probabilistic interpretations for MYCIN's certainty factors, *Proc. Workshop on Certainty and Probability in Artificial Intelligence*, UCLA, Los Angeles (1985) 9-20.

- [24] M. Ishizuka, K. S. Fu and J. T. P. Yao, A rule-based inference with fuzzy set for structural damage assessment, in: M. Gupta and E. Sanchez, Eds., Fuzzy Information and Decision Processes (North-Holland, Amsterdam, 1982).
- [25] A. Kaufmann and M. M. Gupta, Introduction to Fuzzy Arithmetic (Van Nostrand Reinhold, New York, 1985).
- [26] C. A. Kulikowski, AI methods and systems for medical consultation, IEEE Trans. Pattern Anal. Machine Intell. (1980) 464-476.
- [27] E. H. Mamdani and B. R. Gaines, Fuzzy Reasoning and its Applications (Academic Press, London, 1981).
- [28] J. McCarthy, Circumscription: A non-monotonic inference rule, Artificial Intelligence 13 (1980) 27-40.
- [29] R. Michalski, J. G. Carbonell and T. M. Mitchell, Machine Learning (Tioga Press, Palo Alto, 1983).
- [30] M. Mizumoto, S. Fukami and K. Tanaka, Fuzzy reasoning methods by Zadeh and Mamdani, and improved methods, Proc. Third Workshop on Fuzzy Reasoning, Queen Mary College, London (1979).
- [31] D. Nau, Expert computer systems, IEEE Computer 16 (1983) 63-85.
- [32] C. V. Negoita, Expert Systems and Fuzzy Systems (The Benjamin/Cummings Publishing Co., Menlo Park, 1985).
- [33] N. Nilsson, Probabilistic logic, SRI Tech. Note 321, Menlo Park, CA (1984).
- [34] A. I. Orlov, Problems of Optimization and Fuzzy Variables (Znaniye, Moscow, 1980).
- [35] H. Prade, Lipski's approach to incomplete information data bases restated and generalized in the setting of Zadeh's possibility theory, Information Systems 9 (1984) 27-42.
- [36] R. Reiter and G. Criscuolo, Some representational issues in default reasoning, Computers and Mathematics 9 (1983) 15-28.
- [37] N. Rescher, Plausible Reasoning (Van Gorcum, Amsterdam, 1976).
- [38] E. Sanchez, Medical diagnosis and composite fuzzy relations, in: M. M. Gupta, R. K. Ragade and R. Yager, Eds., Advances in Fuzzy Set Theory and Applications (North-Holland, Amsterdam, 1979) 437-444.
- [39] G. Shafer, A Mathematical Theory of Evidence (Princeton University Press, Princeton, 1976).
- [40] E. H. Shortliffe and B. Buchanan, A model of inexact reasoning in medicine, Mathematical Biosciences 23 (1975) 351-379.

- [41] P. Szolovits and S. G. Pauker, Categorical and probabilistic reasoning in medical diagnosis, *Artificial Intelligence* 11 (1978) 115-144.
- [42] L. Wesley, J. Lowrance and T. Garvey, Reasoning about control: an evidential approach, SRI International Tech. Rep. 324, Menlo Park, CA (1984).
- [43] R. R. Yager, Quantified propositions in a linguistic logic, in: E. P. Klement, Ed., *Proc. 2nd Int. Seminar on Fuzzy Set Theory*, Johannes Kepler University, Linz, Austria (1980).
- [44] L. A. Zadeh, Outline of a new approach to the analysis of complex systems and decision processes, *IEEE Trans. Systems Man Cybernet.* 3 (1973) 28-44.
- [45] L. A. Zadeh, The concept of a linguistic variable and its application to approximate reasoning, *Inform. Sci.* 8 (1975) 199-249; 301-357; 9, 43-80.
- [46] L. A. Zadeh, PRUF--a meaning representation language for natural languages, *Internat. J. Man-Machine Studies* 10 (1978) 395-460.
- [47] L. A. Zadeh, Fuzzy sets and information granularity, in: M. Gupta, R. Ragade and R. Yager, Eds., *Advances in Fuzzy Set Theory and Applications* (North-Holland, Amsterdam, 1979) 3-18.
- [48] L. A. Zadeh, A theory of approximate reasoning, *Electronics Research Laboratory Memorandum M77/58*, University of California, Berkeley (1977). Also in: J. E. Hayes, D. Michie and L. I. Kulich, Eds., *Machine Intelligence* 9 (Wiley, New York, 1979) 149-194.
- [49] L. A. Zadeh, Fuzzy sets as a basis for a theory of possibility, *Fuzzy Sets and Systems* 1 (1978) 3-28.
- [50] L. A. Zadeh, Inference in fuzzy logic, *Proc. 1980 ISMUL*, Northwestern University, Evanston (1980).
- [51] L. A. Zadeh, Test-score semantics for natural languages and meaning-representation via PRUF, *Tech. Note 247*, AI Center, SRI International, Menlo Park, CA (1981). Also in: B. B. Rieger, Ed., *Empirical Semantics* (Brockmeyer, Bochum, 1981) 281-349.
- [52] L. A. Zadeh, A computational approach to fuzzy quantifiers in natural languages, *Computers and Mathematics* 9 (1983) 149-184.
- [53] L. A. Zadeh, Formalization of common sense reasoning based on fuzzy logic, *Proc. AIAA*, Long Beach (1985)
- [54] L. A. Zadeh, Syllogistic reasoning in fuzzy logic and its application to usuality and reasoning with dispositions, *IEEE Trans. on Systems, Man and Cybernetics* (1985).

- [55] M. Zemankova-Leech and A. Kandel, Fuzzy Relational Databases--A Key to Expert Systems, Interdisciplinary Systems Research, Verlag TUV Rhineland, Cologne (1984).
- [56] A. Zimmer, Some experiments concerning the fuzzy meaning of logical quantifiers, in: L. Troncoli, Ed., General Surveys of Systems Methodology (Society for General Systems Research, Louisville, 1982) 435-441.
- [57] H.-J. Zimmermann and P. Zysno, Latent connectives in human decision making, Fuzzy Sets and Systems 4 (1980) 37-52.

New Tools for Emergency Management

David L. Waltz

Thinking Machines Corporation
Cambridge, MA

and

Brandeis University
Waltham, MA

ABSTRACT

This paper explores a number of possible applications of AI (Artificial Intelligence) to the problems of emergency management. First, it presents some far-future scenarios, illustrating the kinds of aid that AI could ultimately supply. While many of the applications will require many years of research before they can be realized, there are some near-term applications, especially for new natural language processing technologies; a few interesting possibilities are discussed. There are marked similarities between the problems of emergency management and the three key problems being addressed by the U.S. Department of Defense; these similarities are identified, and projections made about "trickle-down" effects of this DoD research. Finally, the paper discusses some obstacles to achieving the goals of automated emergency management systems and aids, and presents some ideas on how we can achieve these goals.

1. Far Future Emergency Management Scenarios

Projecting current research in AI, we can envision emergency management methods of the future that are dramatically different from today's practice. For example, consider the following possible scenarios (also see Figure 1):

Forest Fire: an automated monitoring system, using satellites and a network of ground-based TV cameras networked with a state-wide emergency management coordination system, catches a forest fire in its early stages. Responsible state officials are called at home and notified of the fire by the system, and the system informs the governor of the plan that it has devised, including transport of firefighting equipment and support supplies from cities A, B, and C to the site of the fire. The governor gives the OK, and the system begins calling the transport drivers to report to work. It issues orders to the robot firefighters, actually bulldozers equipped with vision and on-board intelligence, built to withstand very high temperatures, some equipped with tanks and nozzles for water or chemicals, others equipped with blades and cutters for clearing forest firelines, still others equipped for refueling or servicing. The robot firefighters begin loading themselves onto the appropriate trucks to await the arrival of their drivers.

Drone aircraft are automatically called by the system and given instructions. They begin runs over the fire, dropping water and chemicals well before the bulldozers arrive. They drop their loads on spots picked out in part by the overall system (based on its assessment of the weather, direction of the fire spread, endangered buildings, natural firebreaks, etc.) and in part on the on-board intelligence of each aircraft, which pinpoints the most appropriate exact point to drop its load so that it will have the maximum retardant effect.

The central system periodically informs the governor of progress. Meanwhile, electronic fund transfers are begun to pay for fuel and supplies used, and orders are placed to replenish stocks of spare parts. Any planes or bulldozers lost are reordered. (They are quite inexpensive: they are built in totally automated factories, do not have to

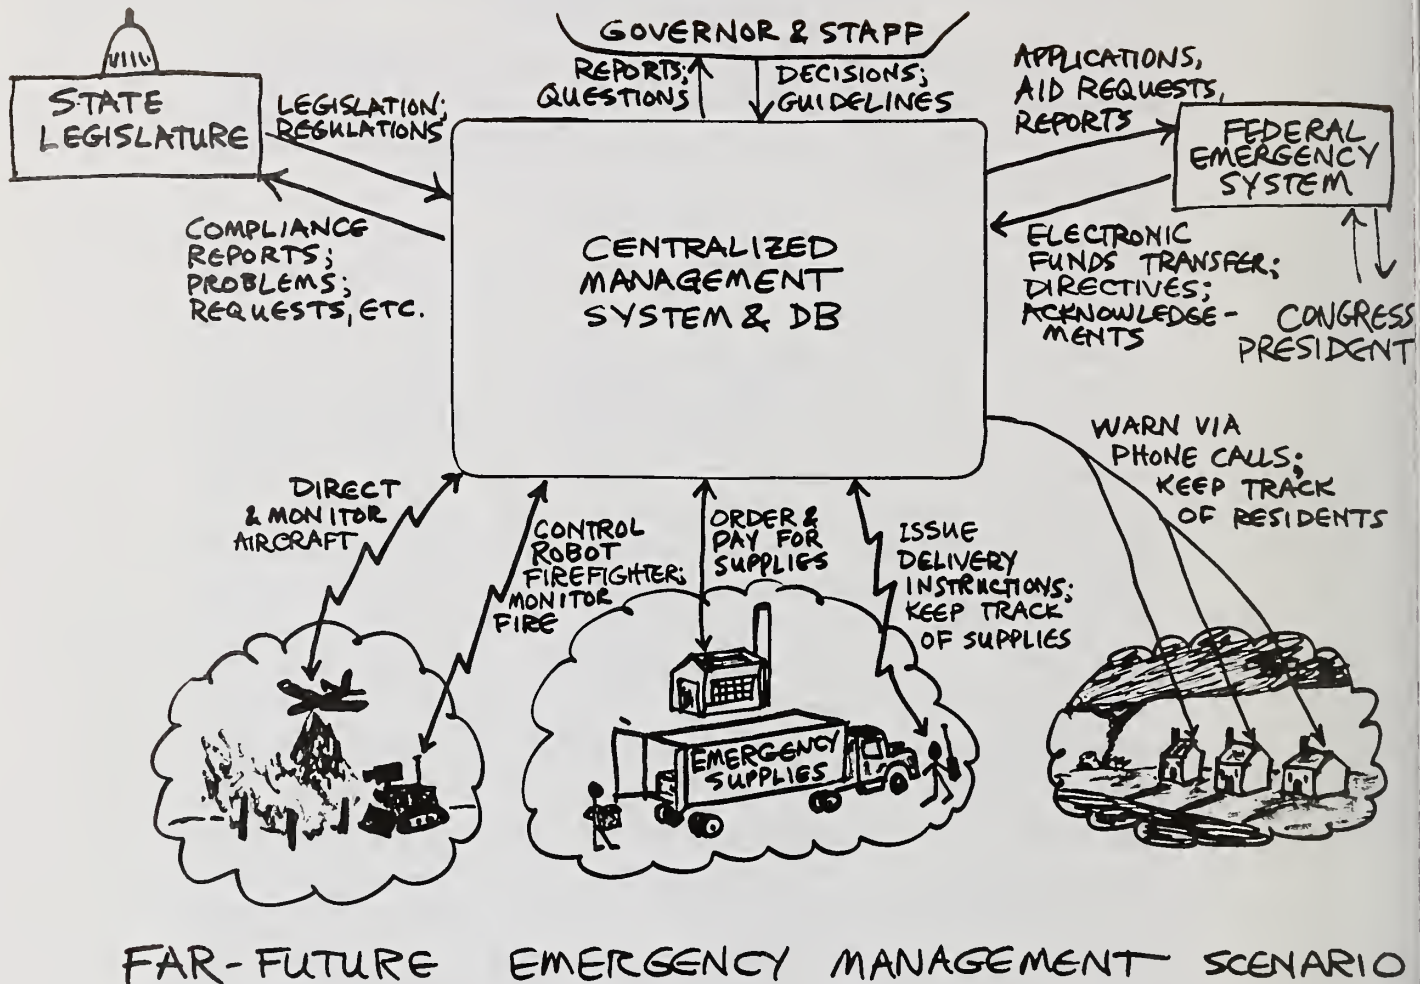


Figure 1

make provision for protecting human operators, and the several powerful computers on board add only a small amount to the total cost of the vehicles.) As the fire is brought under control, another fire breaks out in another part of the state, and some of the vehicles are instructed to go to the new site. They have been working 24 hours a day for three days; the robots firefighters that have sustained damage "volunteer" to go to the other site, and are automatically repaired enroute. Since they were designed to be repaired automatically, they have on-board self-diagnosis facilities, and repair robots can remove and replace most components as entire modules to effect repairs.

Tornado: When a tornado is sighted by a distributed computer vision system, the automated centralized regional emergency system (containing map, resident, and phone number data) calls the residents in and near its path, causing a special emergency ring on their phones, and interrupting calls in progress. Following the tornado, the robot rescue and clean-up cousins of the firefighters are deployed to the scene of the damage. Meanwhile, the system prepares detailed applications for federal aid, based on an automated damage assessment system that gets its data from the robots at the scene, as well as from real estate value data bases. The same system will later file follow-up reports automatically. As regulations change, federal computers transmit program alterations to state and regional systems, and the changes are reflected in the operation of these systems with little or no need for human intervention. When it becomes clear that the current number of emergency vehicles was inadequate for dealing with the tornado emergency, the system helps to generate a system upgrade request to the state legislature, providing detailed justification in terms of the savings in lives, property, and recovery that could have been realized if the requested equipment had been available. This justification is backed up by a computer simulation of the event, illustrating the effects of a variety of possible changes in the emergency aid delivery system, and incorporating long-term weather norms to assess expected costs and benefits.

Other emergencies: The same technologies could generate robot emergency handlers for dam repair; bomb defusing; mine rescue (or

mining itself, eliminating mine rescue); cleanup of railroad or truck accidents involving toxic chemicals; cleanup of hazardous waste dumps; repair or cleanup of nuclear reactors; early warning, evacuation and rescue, and cleanup from earthquakes and floods; and so on. Skill and intelligence could be brought to bear on problem situations without endangering the lives of the rescuers. In many situations, this might allow us to avoid emergencies: for example, the quick righting of a derailed tank car by a fearless, strong, expendable robot might prevent the need for evacuation; early personalized warning of tornados or floods could save many lives; and so on.

2. What problems need to be solved to achieve this future?

The systems described above are large, costly, and very complex by today's standards. Many components of the system are little more than small-scale laboratory experiments today. However, there is good reason to believe that this future is possible, perhaps even likely. Some investment in and encouragement of the dream now could help bring this future about sooner.

First, consider some of the technical advances that need to be made:

- (a) Centralized knowledge bases must be built up, along with mechanisms to keep them up-to-date. Such knowledge bases must contain information on residents within a jurisdiction, phone numbers, map and route data; emergency equipment and its location, condition, abilities, support needs; local, regional, state and federal authorities responsible for making and carrying out decisions; histories of past emergencies for aid in understanding needs and strategies for dealing with future ones in a better manner; rules, regulations, reporting and request procedures and contacts, and so on.
- (b) Fast, cheap, reliable, and smart sensors (optical, seismic, thermal) must be devised and deployed.

- (c) A vast communications system must be constructed. It must be capable of high data-rate transmission and sufficiently redundant or rugged to operate continuously, even under extreme emergency conditions.
- (d) Intelligent systems must be devised that can:
- analyze, sort, and integrate incoming messages and data from people, robots, distributed sensors, and satellites, to form an accurate assessment of any given situation.
 - form plans for allocating resources, deploying equipment, informing the appropriate persons (potential victims as well as decision-makers), gathering data needed to comprehend a situation, and keeping costs as low as possible.
 - give detailed direction to the actual emergency handlers, human and robot, revising directions as each situation changes.
 - formulate plans for improving preparedness, based on advances in technology and shifts in demographics, and help prepare requests and justifications for these plans.
 - prepare emergency aid requests, submit aid applications, help in distributing aid, file follow-up reports, suggest procedure revisions, keep histories.
- (e) Robot emergency handlers must be developed. These systems must have
- good sensory abilities;
 - the ability to take general instructions and adapt them to an actual situation;
 - the ability to construct and send accurate, appropriate reports to the central emergency handling system;
 - and, of course, the manual or physical abilities and dexterity to handle particular emergencies.

Some of the special technical complications involved include:

- the inherent unpredictability of emergencies; the great diversity of types of data and communications systems that must be integrated;

- the problems of making assessments from incomplete or conflicting data;
- the very large amounts of computing power that must be deployed at many levels: small, powerful, rugged, portable units for robots; powerful, reliable units for distributed sensors; and ultrapowerful units for centralized data integration, situation assessment, communications handling, plan generation, and detailed direction of emergency handlers;
- the abilities of all systems to withstand extremes of temperature and lighting, precipitation, power interruption, the shocks of earthquakes and explosions; and so on.

Some other special problems include:

- the planning for and gradual integration of such systems into our existing infrastructure. Some acceptable balance must be struck between completely autonomous systems, and systems which offer possible courses of action for human approval – each type of system may have its own range of applicability.
- the sensible assignment of responsibility for various types of emergencies in various locations – there are problems both with jurisdictional gaps that aren't covered at all, and with overlaps of jurisdiction where systems may work at cross-purposes;
- and, of course, the cost of constructing such systems, which will clearly be very high, even if much of the hardware needed continues its steep long-term historical price decline.

3. What reason do we have to suppose that this future is possible?

Virtually every one of the problems that needs to be solved is a current active research area, though the applications envisioned for the research are for the most part completely outside the domain of emergency management.

In particular, in preparing this paper, I was struck by the very close similarities between the types of problems that need to be solved to realize the far-future emergency scenarios, and the three problems that form the core of DoD's current AI research effort. The three DoD problems are:

- Battlefield management;
- Autonomous vehicles; and
- Pilots' associates.

Of these, the first two are the most relevant. The "Battlefield Manager", like the emergency manager, is expected to integrate data from remote sensors, voice reports, etc., and form some overall assessment of a threatening situation. Both systems must plan, deploy and direct humans and robots. In both cases, response speed is critical. Both must deal with incomplete and conflicting information: in the emergency case, we assume that nature is capriciously unpredictable, while in the battlefield, an enemy may deliberately try to mislead and confuse the system.

The "Autonomous Vehicle" is a robot tank, reconnaissance, or supply-carrying vehicle. Many of the problems for an emergency-handling vehicle are analogous to the DoD autonomous vehicle: both must operate in all types of weather and terrain, and under adverse additional conditions: hostile enemy action in the DoD case; fire, flood, toxic chemicals, landslide, earthquake, explosion, and so on in the emergency case.

While less relevant, research on a pilot's associate might be a help in developing drone emergency handling aircraft, and certainly, civilian pilots might also benefit in avoiding or handling emergencies from having pilots' assistants.

In addition, much of the commercial work on expert systems for handling electronic communication, financial planning and transaction processing, decision-making, and the simulation of large-scale systems will be of benefit in the emergency management realm, as will

research on industrial automation and robotics, signal processing, and computer vision.

4. What can be done now?

Some current research projects and commercial products can already be of aid to emergency managers. I will mention a few examples in the areas of expert systems and natural language processing. For a much more detailed assessment of these two areas see [Waltz 1983a] and [OTA 1985].¹

Expert systems: most expert systems built to date have 1000 rules or less. Nearly all successful applications have been for well-circumscribed, fairly well-understood areas. Such systems are already quite hard to devise; generally a human expert and one or more programmers must spend a substantial amount of time (on the order of a year) to build one, though there are a number of expert system-building tools on the market that promise to cut that time, as well as ease the problems of getting a system to the point where one can begin to trust its reliability. There is general agreement that the bottleneck in developing expert systems is the problem of knowledge acquisition [Davis 1982]. A fair amount of work has gone into trying to speed up expert systems; for an assessment of the limits of speed-ups, see [Forgy et al 1984].

Natural language: a number of tools now exist, ranging from text preparation aids (for checking spelling, grammar, and style of text); to data base question-answering systems that allow a user to phrase questions in English [Waltz 1978, Hendrix 1978, Harris 1977, Woods et al 1972]; to automatic indexing systems for helping to organize documents and messages, and for helping to generate indexes for proposals or reports. Again there are tools available for building custom NL applications.²

¹The November 1983 issue of the IEEE Spectrum also is a good source for information on knowledge acquisition [Hayes-Roth 1983], natural language processing [Waltz 1983b], speech understanding [Reddy 1983], and vision [Kanade 1983].

²For example, the IRUS system from Bolt Beranek and Newman, Inc., Cambridge, MA, or the KLAUS system from SRI International, Menlo Park, CA.

Table I shows a very abbreviated summary history of natural language projects, both laboratory experiments and applications.

In the next section, I give a little more detail on what natural language systems can offer to emergency management systems in the near and far term. In part this is because natural language systems are my own interest, and the area I know best; while they are also quite relevant, I will not treat expert systems in much more detail, because other papers in this volume are already devoted to them.

5. Role of natural language processing systems

Natural language processing (NLP) systems can be useful because:

- any database items are in natural language (NL) already;
- Emergency handlers and the general public already know NL (e.g. English);
- Existing NL tools can provide immediate help for a variety of problems;
- Experience with interim NL systems can help in developing ultimate expert systems, by helping to build up the knowledge bases that will be needed.

Figure 2 shows the places where NL will be needed in a full-scale emergency management system. NL is needed to interpret reports phoned or typed into the system from the field; to route this information to the right parts of the database; to handle questions in NL posed by users of the system; to process regulations and fill in report forms (while following regulations); to prepare proposals and requests for state and federal aid; and so on.

The TMC Indexer³ is an example of a product that can be useful right away. It is a system that takes text, and automatically generates a back-of-the-book like index for the text. It can be used to keep related topics together: it can organize messages, proposals, reports, plans,

³The TMC Indexer is available from Thinking Machines Corp., Cambridge, MA.



APPLICATION

- * **Experimental machine translation**

1960

- * AI: question-answering; problem-solving; robot control

- * DB keyword search
- * Word processing systems

1970

DB Q/A Systems LUNAR (Woods)
LIFER (Hendrix)
PLANES (Waltz)

- * **Story understanding (Schank)**
- * **FRUMP (DeJong)**
 - News story summarizer

- * **Spelling correctors**
- * **DB question-answering**
(Harris: INTELLECT, AI Corp.)
- * **Isolated-word speech recognition**

1980

- * NL Understanding
 - Stories
 - Dialogues
 - Articles

- * **Message handling (Hardt)**
- * **Expert system explanation generators**
- * **NL menu (Tennant: TI)**
tailorable NL interface

- * Limited machine translation

- ### * NL Generation

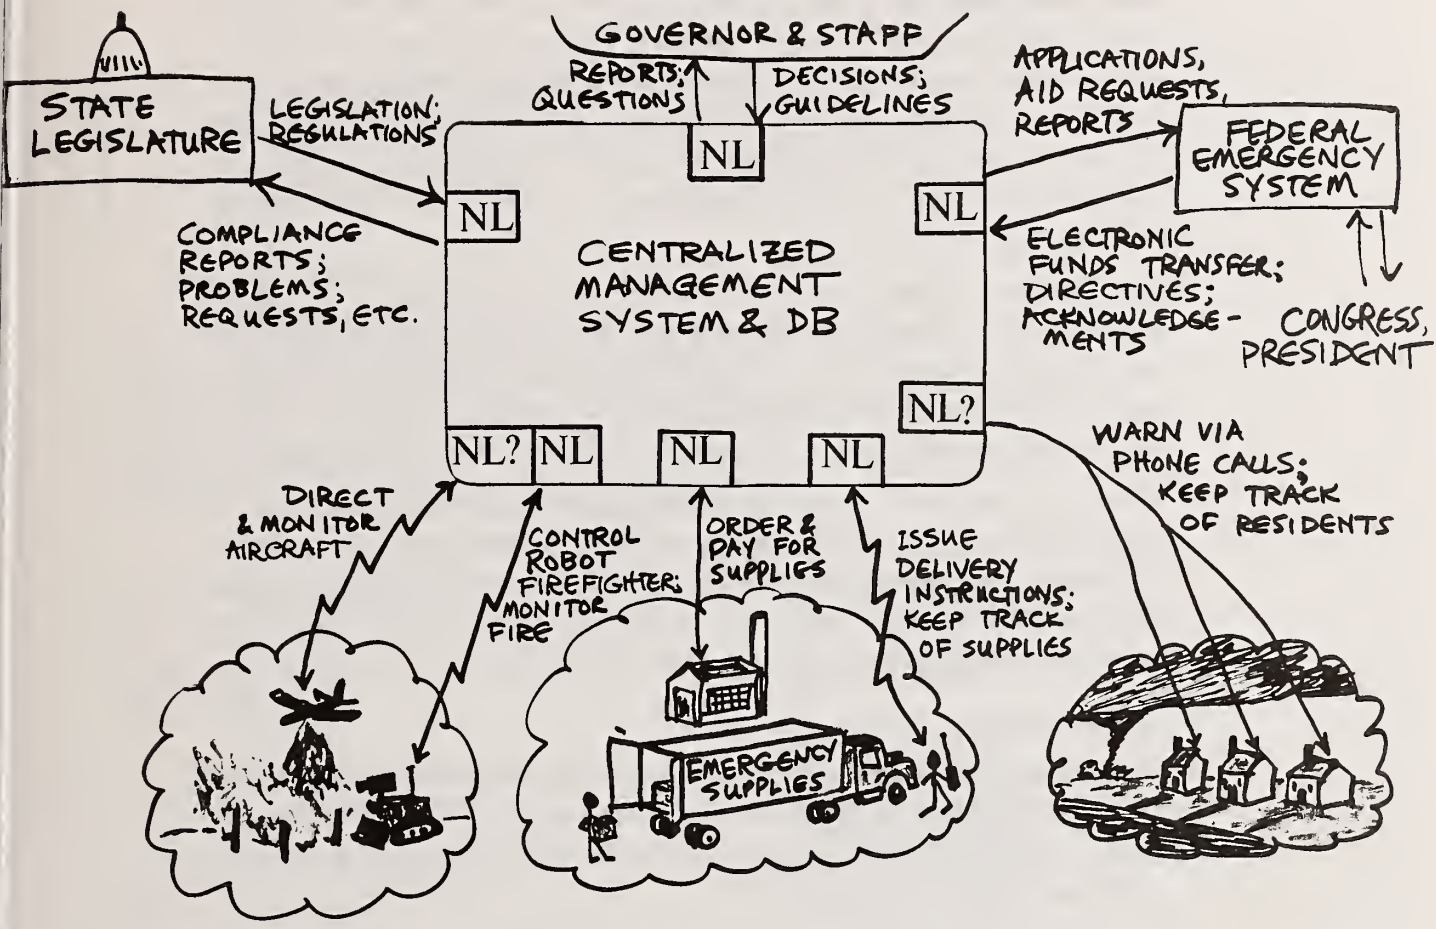
- * **Style checker (Heidorn: IBM)**

1985

- ## * Theory

- * **Automatic subject index generator**
(Mott, Waltz, Resnikoff: TMC Indexer)
- * **On-line help systems**

Table 1



FAR-FUTURE EMERGENCY MANAGEMENT SCENARIO

Figure 2

histories, etc. by finding key terms, and generating an alphabetic index for any or all of these items. The Indexer contains a handy interface that allows a user to see the index in one screen window, and the original text in the other window. When the user picks a term in the index (using either a mouse or a key) the text in which the term occurs pops up in the text window, and the term is highlighted in both windows, as illustrated in figure 3. (See [TMC 1985]) An Indexer would be a useful component at each of the sites marked "NL" in figure 2.

The Intellect system ⁴ is a commercially available NL database question answering system; it allows a user to ask questions in relatively free English [Harris 1977]. English language database "front ends" can be very useful for providing casual or infrequent users with access to information in a data base.

When coupled with a speech understanding system of the sort expected to become widely available within the next ten years, and an automatic document reader (such as the Kurzweil reader), systems like the TMC Indexer and Intellect can form a potent, useful core for emergency management aid systems of the future.

6. Other issues

There are political and ethical problems with constructing systems of the sort described in the far-future scenarios. I will just mention two here, though I suspect many more issues will surface over the next few years.

- Ethical considerations: in some circumstances, a system may have to decide to meet one of two or more simultaneous emergencies first, because of resource limitations. The system will thus potentially make decisions to let one group suffer or die in order to save another. In a similar case, a poorly designed system might fail to prevent or defuse an emergency. In such cases, who is responsible? The designers of the system? The operators of the system? The agency that contracts for its design?

⁴The Intellect system is available from AI Corp., Waltham, MA.

Text File: INDEXER: DOC; SAMPLE.TEXT (NORMAL, 6156 bytes, 121 lines)
 Index File: indexer:doc;sample.text (85 entries from 1 files) [Read-Only]

TMC Indexer 1.8

[Top of Text](#)

Description of the Connection Machine Computer

Dr. David L. Waltz

1. The Hillis Connection Machine Concept

The Connection Machine is a computer for concurrently manipulating large amounts of information, which may be stored in any convenient form. For a broad range of problems, the Connection Machine provides processing power proportional to the size of a problem. Each data element or structure of data elements can be assigned its own simple processor in the Connection Machine, up to perhaps a million processor/memory cells. Each processor can be dynamically connected to any other processor, so that the geometry of a problem can be mirrored by the interconnection structure of the processors. If the interconnection structure consisted of physical wires, the machine would have to be rewired for every problem. Since this is impractical, the processor cells are connected through a Packet Switched Network. They communicate by sending messages. Receiving a message causes a

[More below](#)
[Top of Index](#)

Advantages , Connection Machine	85Z (inversions: 1 below)
advantages , Connection Machine offers	74Z (inversions: 1 below)
AI	51Z
area , chip	87Z
Artificial Intelligence , field of	58Z (inversions: 2 below)
Artificial Intelligence , problems in	62Z (inversions: 2 below)
Central Processing Unit	22Z
Codd's relational algebra	92Z
computer , serial	38Z (inversions: 1 below)
	48Z (inversions: 1 below)
computers , serial	36Z (inversions: 1 below)
	45Z (inversions: 1 below)
	63Z (inversions: 1 below)
Computing Needs	48Z
Connection Machine	2Z (inversions: 1 below)
	4Z (inversions: 1 below)
	17Z (inversions: 1 below)
Connection Machine , features of the	86Z (inversions: 2 below)
Connection Machine , processor in the	7Z (inversions: 2 below)
Connection Machine Advantages	85Z (inversions: 1 above)

[More below](#)

84/14785 11:02:46 Janet

USER:

lvi

DFILE serving BUTLPH

Figure 3

Who will decide the policies to be built into such systems? How will they weight the costs and benefits of various courses of action, e.g. what are the relative values of life and property? Will the public take part in the design of policies for such systems? Should it?

- There are real dangers to centralized systems that know a great deal about citizens, and have distributed information- gathering potential; they have potential for misuse as "Big Brother"; there may be political resistance to deploying such systems.

It is not too early to consider and prepare to deal with such problems. In fact, Isaac Asimov, George Orwell, Robert Heinlein, and other science fiction writers have considered such issues in considerable depth over the last 40 years.

However, we should not lose sight of the fact that such systems can save lives and property that are now lost, without risking the lives of rescuers. Such systems should in the long run make us all relatively free of fear and danger from natural and man-made catastrophes.

References

[Davis 1982] Davis, R. and D. Lenat. *Knowledge-Based Systems in Artificial Intelligence*. New York: McGraw-Hill, 1982.

[DoD 1983] Strategic Computing, "New generation computing technology: a strategic plan for its development and application to critical problems in defense", Defense Advanced Research Projects Agency, Oct. 28, 1983, p. vii (citation from [OTA 1985]).

[Forgy 1984] Forgy, C., A. Gupta, A. Newell, and R. Wedig. "Initial Assessment of Architectures for Production Systems" *Proc. AAAI-84 (1984 National Conference on Artificial Intelligence)* Austin, TX, August 1984, 116-120.

[Harris 1977] Harris, L. "User oriented data base query with the ROBOT natural language query system", *International Journal of Man-Machine Studies* 9, 697-713.

[Hayes-Roth 1983] Hayes-Roth, F. "Codifying knowledge for machine reading" *IEEE Spectrum* 20, 11, 79-81, November 1983.

[Hendrix 1978] Hendrix, G. G., E. D. Sacerdoti, D. Sagalowicz, and J. Slocum. "Developing a natural language interface to complex data" *ACM Transactions on Database Systems* 3, 2.

[Kanade 1983] Kanade, T., and R. Reddy, "Computer vision: the challenge of imperfect inputs" *IEEE Spectrum* 20, 11, 88-91, November 1983.

[OTA 1985] *Information Technology R & D: Critical Trends and Issues*. Washington, DC: U.S. Congress, Office of Technology Assessment, OTA-CIT-268, February 1985.

[Reddy 1983] Reddy, R. and V. Zue, "Recognizing continuous speech remains an elusive goal" *IEEE Spectrum* 20, 11, 84-7, November 1983.

[TMC 1985] *The TMC Indexer Manual*. Thinking Machines Corp., Cambridge, MA.

[Waltz 1978] Waltz, D. L. "An English language question answering system for a large relational database" *Communications of the ACM* 21, 7, 526-39.

[Waltz 1981] Waltz, D. L. "The State of the Art in Natural-Language Understanding". in W. Lehnert and M. Ringle (eds.) *Strategies for Natural Language Processing*, Hillsdale, NJ: Lawrence Erlbaum Associates, 3-32.

[Waltz 1983a] Waltz, D. L., M. Genesreth, P. Hart, G. Hendrix, A. Joshi, J. McDermott, T. Mitchell, N. Nilsson, R. Wilensky, and W. Woods. "Artificial Intelligence: An assessment of the state-of-the-art and recommendation for future directions" *AI Magazine* 4, 3, 55-67, Fall 1983.

[Waltz 1983b] Waltz, D. L., "Helping computers understand natural languages" *IEEE Spectrum* 20, 11, 81-4, November 1983.

[Woods 1972] Woods, W. A., R. M. Kaplan, and B. Nash-Webber. 1972. *The lunar sciences natural language information system: final report*. Report No. 2378, Bolt Beranek and Newman, Inc., Cambridge, MA.

TIMM™ - A TOOL FOR BUILDING EXPERT SYSTEMS

Joseph N. Craig
General Research Corporation
7655 Old Springhouse Road
McLean, Virginia 22102

INTRODUCTION

TIMM™, the Intelligent Machine Model, is a computer program that is used to build expert systems.

Feigenbaum [1] has defined an expert system as follows:

"An expert system is an intelligent computer program that uses knowledge and inference procedures to solve problems that are difficult enough to require significant human expertise for their solution. The knowledge necessary to perform at such a level, plus the inference procedures used, can be thought of as a model of the best practitioners in the field.

The knowledge in expert systems consists of facts and heuristics. The 'facts' constitute a body of information that is widely shared, publicly available, and generally agreed upon by experts in a field. The 'heuristics' are mostly private, little-discussed rules of good judgment (rules of plausible reasoning, rules of good guessing) that characterize expert-level decision-making in the field. The performance level of an expert system is primarily a function of the size and quality of the knowledge that it possesses."

Typical expert systems consist of two primary components.

- (1) a knowledge base of domain facts and heuristics associated with the problem, and
- (2) an inference procedure (or control structure) for using the knowledge base in the solution of the problem.

Frequently expert systems are augmented by a variety of "user friendly" functions that ease communication with and understanding of the expert systems. Additions include "natural language" interfaces and "explanation facilities." The latter permit a user to challenge or examine the reasoning process that leads to the system's decision.

The development of domain specific knowledge bases is the key to expert systems technology. In fact, an entire field, "knowledge engineering," has developed. Knowledge engineers are artificial intelligence and computer specialists who interview experts and translate their expertise into the set of "rules" that constitute a knowledge base.

Because knowledge engineers are a scarce and expensive resource, the development of expert systems has tended to be an expensive proposition.

Recently, several products have appeared on the market which provide software environments in which expert systems can be developed and used. TIMM™ is such a product. TIMM™ provides all of the tools needed by an expert to define, train, and refine an expert system. TIMM™ also provides an environment in which a non-expert can use expert systems, ask questions, and obtain expert-level advice.

TIMM™ has advantages over other products that make it attractive. First, TIMM™ is highly portable and is available on a wide variety of standard computer systems. TIMM™ is written in FORTRAN so it can be interfaced or embedded in other programs relatively easily. Second, TIMM™ provides an English-like interface permitting experts to develop expert systems directly. Tools available within TIMM™ allow the expert to define an expert decision structures, provide rules to the system, customize the user interface, and iterate and modify the system until acceptable performance has been achieved. All of these can be accomplished without any programming effort. Third, TIMM™ employs a unique analogical inferencing procedure to reach decisions. Finally, a multiple expert system can be linked together.

2. TIMM™ PHILOSOPHY

Automated, rapid acquisition and representation of knowledge is a desirable goal for expert system development. In the development of TIMM™, three problem areas that exist in the traditional approach to expert system development have been addressed specifically. These problems are: the time and cost to develop knowledge bases is high; the constructive aspect of problem solving is often hidden; and reconciling the continuously varying nature of our experience with the exact logical modeling of standard inference schemes is difficult. The first of these problems is a practical one. The others, involving constructive and analogical thought, directly apply to the modeling of human problem solving.

We describe these problems below and show how TIMM™ deals with them.

The Time and Cost to Develop Knowledge Bases is High

Knowledge acquisition and representation is hard and slow [2, 3, 4]. Davis has observed that "... even for the best-understood problems, experienced researchers using the best-understood technologies still require at least five man-years to develop a system that begins to be robust" [5].

The basis of this problem is the need for specific world knowledge for intelligent behavior. It is now recognized that the power of expert systems comes from large amounts of knowledge [5, 6]. Two needs of knowledge acquisition stand out. First, a way is needed to acquire and represent knowledge in a way congenial to human understanding (as distinguished from representational convenience). Second, we need to automate the process of knowledge acquisition, to simplify and accelerate the process and to remove the need for specialized engineering skills to achieve usable knowledge representation.

Construction in Problem Solution Is Hidden

Most expert systems approach problem solving through the repetitive reading of lists of rules, terminating when a satisfactory candidate solution is isolated. It is essentially a process of elimination. But human problem solving proceeds by constructing and modifying models [7, 8, 9]. Structure is important in human problem solving.

- Meaning guides convergence on problem solutions. But meaning is a product of context, which, in the importance of order in creating meaning, entails structure [10].
- Aesthetics in problem solution allows construction of elegant solutions. Elegance is defined by what is known of past solutions, compared with whatever partial solution is known to exist for the current situation. This is a structural relation [7].

- Analogies, metaphors, cues, paradigms, and expectations are used in constructing solutions and explanations. All of these involve problem structures.

In most expert systems, structure and construction are not modeled explicitly [11]. The A.I. community has become more aware of the importance of separating declarative and procedural knowledge [12], and frame systems have been developed to reconcile the need for cooperation between declarative and procedural knowledge in mapping the structures of a problem domain [13, 14, 15]. The fluidity of human problem solving has not however been captured.

Continuously Varying Experience Doesn't Fit Exact Inferencing

The importance of explanation is impressing itself upon workers in expert systems [10, 16]. The psychological underpinnings of expert systems should be such that explanations flow from a problem solving model compatible with that of the system's users.

Exact match inferencing may be too rigid a model for human problem solving [17]. Work has been done on partial match rules to allow recognition of incomplete matching. Still, partial match inferencing is organized around exact matching of partial groupings of decision criteria, rather than on inexact or analogous matching.

People use knowledge of continuously varying natural phenomena in problem solving. The world of our perceptions -- a world of sizes, colors, sounds -- is one of continuous variation. It is sensible to talk of "like", "almost", "very similar". Knowledge acquisition and representation schemes which fail to take this into account, and inference schemes which are rigid, do not mirror human problem solving.

Rapid Development of Procedural Knowledge

TIMM™ is a computer program that incorporates approaches to many of the ideas expressed above. It is a domain-independent tool for building expert systems.

TIMM™ separates knowledge into declarative and procedural sections, much like frame-based systems which allow attached procedures. Somewhat surprisingly, the frame paradigm, which organizes knowledge in a set of slots (attributes) rather than in IF-THEN rules, has only been minimally applied to expert systems development.* Use of a frame-like approach

*KNOBS is a frame-based system, but variations of the production rule paradigm are much more common in expert systems. In Building Expert Systems [6], there are only three references to frames in the index, each reference being minor.

allows easy characterization of the stable aspects of the domain. TIMM™ maintains separation of control knowledge from domain facts and relations [10], and follows completion of declarative knowledge definition with a decision-oriented approach to procedural knowledge acquisition.

This decision-oriented approach supports rapid development of procedural knowledge, and is based on the simple observation that experts are better at doing than at describing. The stable knowledge about a problem domain in the declarative knowledge section lays the groundwork for the system's synthesis of realistic problem descriptions. Using these, an interactive dialog can be entered with the expert, oriented toward making decisions (doing), rather than explaining how decisions are made (describing). The system, originating questions and monitoring answers, can create rules based on observed expert behavior (which can then be offered to the expert for annotation.) This process greatly speeds the development of rule bases, and a knowledge engineer is not required.

Modeling Constructive Problem Solving

TIMM™ models constructive problem solving in two ways. First, it allows explicit representation of domain paradigms to guide making individual decisions. Second, it models the expectations of the expert in its domain representation.

TIMM™ allows incorporation of expert-defined paradigms in constructing a problem's search space. Both declarative and procedural knowledge representation allow nesting of varying levels of definition, from the very general to the very specific. Further, overlapping of definition is allowed in the knowledge base. In this way, overlapping high-level structures can represent multiple paradigms valid in the problem domain. These can be used to guide search -- the lower level specifics are not instantiated if there is not a good match with the high-level paradigm. Efficiency of search is aided by guiding the search from the earliest opportunity, as applicable paradigms are discovered and inapplicable paradigms eliminated, so that only useful structures are instantiated. Moreover, explanatory power is gained by having the problem domain paradigms explicitly defined and available for inspection by users.

Expert's expectations about normal domain behavior are incorporated by extending the implementation of the frame paradigm to include a measure of the typicality of the values associated with slots in individual frames. This expands the idea of default slot values to let each slot have a graded set of 'expectations', rather than only one. By also providing a measure of salience among slots within a frame, the varying importance of different aspects of a situation can be conveniently modeled. Typicality and saliency increase the flexibility of frame structures. Since expected (i.e. typical) values are specified, an expectation-guided problem solving strategy is possible. By combining the program's expectations with the gradual instantiation of

paradigms, each succeeding step in problem solution is dependent both on the current state of the system's understanding and on the system's expectations. This is a much more satisfactory model of human problem solving than the stimulus/response model most often used in expert systems.

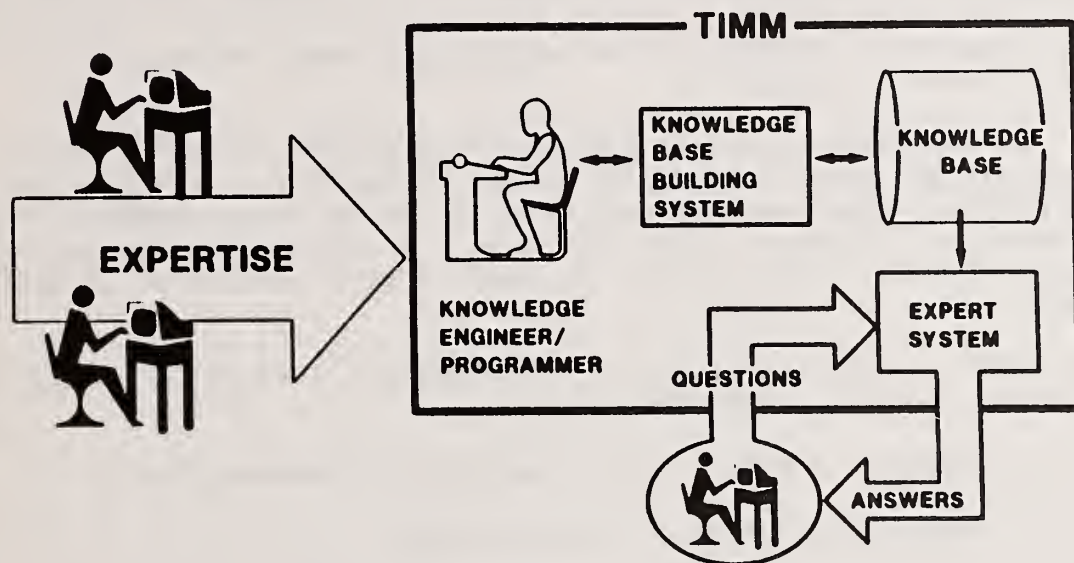
Analogical Inferencing and Knowledge Generalization

TIMM™ approaches the problem of reasoning with incomplete, uncertain, and ambiguous situation knowledge through use of analogical inferencing. A side effect of this technique has been the ability to generalize procedural knowledge.

Reasoning with uncertainty is often necessary. TIMM™ combines partial match inferencing with an ability to measure the similarity of a perceived situation to known situations within individual rule clauses. Partial match inferencing is an important tool for dealing with uncertainty, but as F. Hayes-Roth points out, "... representations should be sought that support the use of proximity measures or directionality to identify good partial matches" [18]. In TIMM™ we have used proximity and directionality -- in general, context-sensitive similarity -- as well as partial matching, in the operation of our inferencing procedure. This approach uses the ranges of values specified for entities in the declarative knowledge section to provide a metric across the problem domain. Similarity of situations can be estimated by their closeness in the representational space. Analogical inferencing is a combination of partial match inferencing with a similarity measure on rule clause matching.

The use of continuously varying knowledge and analogical inferencing allows TIMM™ to generalize procedural knowledge. Because TIMM™ can gauge the similarity between rules, groups of similar rules can be experimentally reduced to single, more general rules. Such rules, if meeting the approval of the expert training the system, can then replace their precursors. While this process sounds simple, some sophistication is necessary to avoid 'stupid' simplifications. In the current version of TIMM™, the salience and typicality of rule clauses and values is implicit rather than explicit. Thus, TIMM™ must examine the content of the rule set to discover which simplifications are based on important similarities, and which are spurious. We hope that the ability to generalize will produce views of the problem domain both unanticipated by, and useful to, the human expert.

A HIGH LEVEL VIEW OF TIMM™



Advanced
Technologies
Division

3. FUNCTIONAL ELEMENTS OF TIMM™

TIMM™ is a domain-independent, knowledge-based program which permits a user to construct, exercise, modify, and interrogate a body of knowledge whose subject matter is selected by the user. TIMM™'s knowledge base contains rules, definitions, etc. about a particular domain. Initially, the knowledge base is empty and an expert adds data to the knowledge base by carrying out a dialog with the "knowledge engineer" that is built into TIMM™. When the knowledge base is developed, a non-expert user can interact through the expert system interface to obtain answers to questions. Both interfaces to TIMM™ are "friendly", providing spelling correction, help facilities, and three levels (terse, verbose, and help) of interaction.

TIMM™ Implementation and Capabilities

The TIMM™ knowledge base is implemented as a data file that is accessed when TIMM™ is exercised. Therefore, multiple experts can be maintained in the computer environment. The user has the option of specifying which expert should be used at a given time. In a multi-user environment, such as VMS on a VAX computer, multiple copies of TIMM™ can run concurrently permitting simultaneous development and/or exercise of multiple expert systems.

The expert system building phase of using TIMM™ centers around dialogue with an expert. First, the expert is led through a definition of the problem. What must be considered in coming to a decision? What decisions are possible? Then, the expert is questioned about specific problem examples. What should be done if this combination of events occurs?

TIMM™ transforms answers to such questions into rules. When the expert is satisfied, TIMM™ can

- eliminate redundant rules (compress the knowledge base),
- discover and ask about inconsistencies,
- say where it doesn't know enough about the problem (check for completeness), and
- look for deeper, simpler rules (generalize), using a pattern recognition procedure).

The problem definition and problem solution dialogues can be reopened, and the process repeated, until the expert is satisfied and system building is complete.

Problems, and our understanding of them, are rarely static. A useful expert system should have the capacity to grow as understanding of the problem grows. TIMM™ makes it easy to add or modify rules when new knowledge comes to light.

The rules generated by TIMM™ are in English. Expert systems built using TIMM™ can explain decisions by displaying the rules used in making the decision. Additionally, the English form of the rules makes inspection of the system's 'expertise' possible by people with no knowledge or interest in programming.



TIMM™ MAIN MENU

(C) COPYRIGHT 1983 GENERAL RESEARCH CORPORATION
TIMM™ THE INTELLIGENT MACHINE MODEL
VERSION 2.0

GOOD MORNING. I'M TIMM™. LET'S BEGIN.

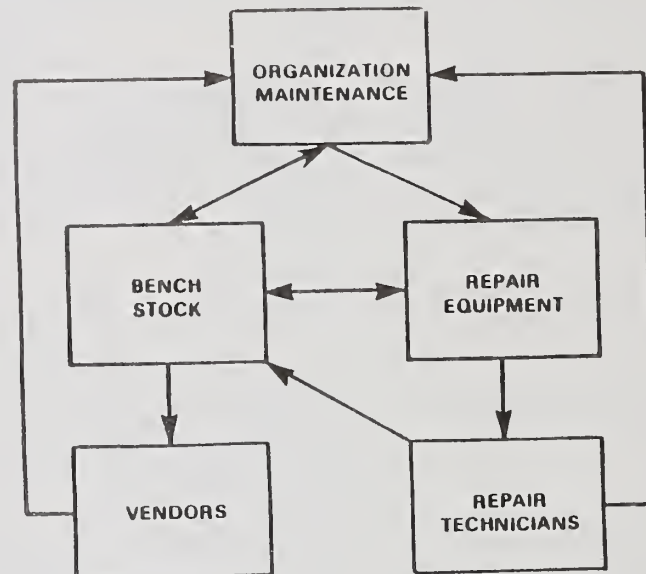
MAIN MENU	:	B	BUILD	T	TRAIN	E	EXERCISE
TYPE LETTER	:	I	INQUIRE	G	GENERALIZE	C	COMPRESS
+ FILE NAME	:	M	MODIFY	N	CHECK CONSISTENCY	P	CHECK COMPLETENESS
AND RETURN	:	R	RECOVER	S	STOP		

SELECTION?

**Advanced
Technologies
Division**

TIMM™ is capable of supporting nine different operations at the user's request. These operations are presented to the user as options from TIMM™'s main menu or as options to functions that can be selected from the main menu.

A LINKED KNOWLEDGE BASE DEMONSTRATION



Advanced
Technologies
Division

To illustrate the use of TIMM™, we present a partial description of the construction of a prototype expert system for "Organization Maintenance." The expert system will help a maintenance activity decide whether to repair a failed part or issue a replacement part from stock. As shown, the system implements expertise in logistics support, inventory control, equipment repair, and vendor selection. Each of these areas is implemented as a separate expert system, linked together using TIMM™'s hierarchical structure.

We focus here on the expert system called "MAINTENANCE ACTIONS." It decides whether to refer to a knowledge base on repair or one on stock management. The figures represent what will appear on the user's terminal screen during a TIMM™ session. User inputs are identified by an arrow (←).



TIMM™ BUILD FUNCTION

(C) COPYRIGHT 1983 GENERAL RESEARCH CORPORATION
TIMM™ THE INTELLIGENT MACHINE MODEL
VERSION 2.0

GOOD MORNING. I'M TIMM™. LET'S BEGIN.

MAIN MENU : B BUILD T TRAIN E EXERCISE
TYPE LETTER : I INQUIRE G GENERALIZE C COMPRESS
+ FILE NAME : M MODIFY N CHECK CONSISTENCY P CHECK COMPLETENESS
AND RETURN : R RECOVER S STOP

SELECTION? B ORGMAIN ←

WHAT WOULD YOU LIKE TO CALL THIS NEW EXPERT SYSTEM: ORGANIZATION MAINTENANCE

WHAT DECISION WILL I BE LEARNING HOW TO MAKE: REPAIR OR REPLACE

TELL ME THE POSSIBLE CHOICES. ENTER ONE PER LINE, AND ENTER A "/" AS THE LAST
LINE:

\$FIX ←
\$BSTOCK ←
/ ←

Advanced
Technologies
Division

3.1 BUILD

TIMM™'s BUILD function guides the user through the process of building a new knowledge base. It first asks the user to specify the file name for the knowledge base and the name of the decision that will be made. Next, the user specifies the possible decisions that can be made. In the example, the two possible decisions are to refer the decision to a repair specialist (FIX) or to an inventory specialist (BSTOCK). The dollar sign (\$) tells TIMM™ that these are the names of other expert systems that should be called when a decision is reached.



TIMM™ BUILD FUNCTION (CONT.)

NOW TELL ME THE NAME OF EACH FACTOR THAT INFLUENCES THE DECISION. ENTER ONE PER LINE, AND ENTER A "/" AS THE LAST LINE:

PART	→
REPAIRABLE	→
PRIORITY	→
REPAIR COST	→
CUMULATIVE REPAIR COST	→
TIME IN SERVICE	→
/	→

PART

ARE THE VALUES OF THIS FACTOR NUMBERS (N) OR DESCRIPTIVE PHRASES (P) ? P →

ARE THESE VALUES UNORDERED (U), LINEARLY-ORDERED (L), OR CIRCULARLY-ORDERED (C)? U →

TELL ME THE POSSIBLE VALUES. ENTER ONE PER LINE, AND ENTER A "/" AS THE LAST LINE:

A	→
B	→
C	→
D	→
E	→
F	→
G	→
H	→
I	→
/	→

Advanced
Technologies
Division

TIMM™ next asks for a list of factors that influence the decision, and for attributes of each factor. The factors can be described as numeric or phrases; and the values can be unordered, linearly ordered, or circularly ordered. If the factor values are unordered, TIMM™'s inferencing procedure will function in a standard forward-chaining manner. If a particular ordering is specified, TIMM™ is able to make use of this knowledge in reaching a decision. [As an example, the days of the week are specified as circularly ordered; TIMM™ would know that Saturday and Monday are equally close to Sunday.]



PLEASE GIVE ME A VERBOSE DESCRIPTION FOR EACH OF THE FACTORS. TERMINATE
EACH DESCRIPTION WITH A "/":

PART

WHICH PART FAILED?/ ←

REPAIRABLE

IS THE FAILED PART REPAIRABLE?/ ←

PLEASE GIVE ME HELP INFORMATION FOR EACH OF THE CHOICES. TERMINATE EACH
ENTRY WITH A "/":

INVOKE FIX

/ ←

INVOKE BSTOCK

/ ←

PLEASE GIVE ME HELP INFORMATION FOR EACH OF THE FACTORS. TERMINATE EACH
ENTRY WITH A "/":

PART

A-AXLE, B-BEARING, C-COUPLER, D-DIAGONAL, E-EAR, F-FLANGE, G-GEAR, I-IBAR. ←

/ ←

REPAIRABLE

IS THE FAILED PART REPAIRABLE LOCALLY?/

← Advanced
Technologies
Division

Help and Verbose prompts are used to provide additional information to the user. When the expert system is being used, the user has the choice of seeing terse or verbose prompts. If HELP is typed at any time, the appropriate help prompt will be displayed.



MAIN MENU : B BUILD T TRAIN E EXERCISE
TYPE LETTER : I INQUIRE G GENERALIZE C COMPRESS
+ FILE NAME : M MODIFY N CHECK CONSISTENCY P CHECK COMPLETENESS
AND RETURN : R RECOVER S STOP

SELECTION? T ←

FOR TRAINING, I NEED A DEFINITE VALUE OR RANGE OF VALUES FOR EVERY FACTOR YOU'VE DEFINED. EITHER YOU PICK THEM, OR I WILL. THEN I'LL ASK YOU FOR YOUR DECISION, GIVEN THOSE VALUES.

SPECIFY VALUES OR RANGES OF VALUES FOR THE FACTORS LIKE THIS:

HEIGHT = 70
WEIGHT = <180 (EQUAL TO OR LESS THAN 180)
RANK = >COL (EQUAL TO OR MORE THAN COL)
SCORE = A:C (BETWEEN A AND C INCLUSIVE)
AGE = * (ANY VALUE -- "IRRELEVANT" OR "DONT CARE")

FOR THE DECISION, YOU CAN SPECIFY JUST ONE CHOICE, OR SPECIFY UP TO TEN CHOICES WITH CERTAINTY NUMBERS (WHOLE NUMBERS P TO 100), LIKE THIS:

PROMOTABLE = YES (67)
= NO (33)
= / (TO INDICATE THAT'S ALL)

SPECIFY THE DECISION AS "#" TO MEAN "THAT COMBINATION OF FACTOR VALUES ISN'T POSSIBLE, OR ISN'T ALLOWED " TO STOP, ENTER A FACTOR VALUE OR CHOICE AS "\$".

**Advanced
Technologies
Division**

3.2 TRAIN

The TRAIN function permits the expert to enter knowledge regarding particular decisions once the domain of the decisions has been defined. TIMM™ provides several ways to enter this information: specific scenarios can be generated by TIMM™, or the expert can enter specific scenarios or more general rules.



SHALL I PICK THE TRAINING CASES FOR YOU (Y OR N)? Y ←

IS IT OK IF I DUPLICATE OLD TRAINING CASES (Y OR N)? N ←

ANOTHER (Y OR N)? Y ←

PART = E
REPAIRABLE = YES
PRIORITY = 2
REPAIR COST = 195
CUMULATIVE REPAIR COST = 237
TIME IN SERVICE = 165
REPAIR OR REPLACE = \$BSTOCK ←

ANOTHER (Y OR N)? Y

PART = A
REPAIRABLE = YES
PRIORITY = 4
REPAIR COST = 145
CUMULATIVE REPAIR COST = 395
TIME IN SERVICE = 15
REPAIR OR REPLACE = \$FIX ←

ANOTHER (Y OR N)? Y ←

Advanced
Technologies
Division

When TIMM™ specifies scenarios, factor values are generated randomly and the expert is prompted for a decision.



SHALL I PICK THE TRAINING CASES FOR YOU (Y OR N)? N →

PART	= A	→
REPAIRABLE	= YES	→
PRIORITY	= 2	→
REPAIR COST	= 500	→
CUMULATIVE REPAIR COST	= 500	→
TIME IN SERVICE	= 30	→
REPAIR OR REPLACE	= \$FIX	

ANOTHER (Y OR N)? Y

PART	= *	→
REPAIRABLE	= NO	→
PRIORITY	= *	→
REPAIR COST	= *	→
CUMULATIVE REPAIR COST	= *	→
TIME IN SERVICE	= *	→
REPAIR OR REPLACE	= \$BSTOCK	→

Advanced
Technologies
Division

When the expert provides examples or rules, TIMM™ prompts for each factor value and the appropriate decision. Here, the second example is a "rule" rather than a specific case. It should be read "If the part is not repairable, then invoke the BSTOCK expert system."


```

*MAIN MENU* :   B BUILD      T TRAIN      E EXERCISE
TYPE LETTER :   I INQUIRE   G GENERALIZE  C COMPRESS
+ FILE NAME :   M MODIFY     N CHECK CONSISTENCY P CHECK COMPLETENESS
AND RETURN :   R RECOVER    S STOP

```

SELECTION? I ←

```

*INQUIRE* :   D SHOW DECISION STRUCTURE  R RETURN TO
TYPE LETTER :   K SHOW KNOWLEDGE BASE      MAIN MENU
AND RETURN :   N SHOW NUMBER OF RULES

```

SELECTION? N ←

SHALL I COPY THE INFORMATION INTO A FILE IN YOUR DIRECTORY AS I DISPLAY IT (Y OR N)? N ←

SIZE OF THE KNOWLEDGE BASE FOR THE EXPERT SYSTEM "ORGANIZATION MAINTENANCE"

```

-----
PRIMARY KNOWLEDGE BASE:    18 RULES
SECONDARY KNOWLEDGE BASE:   0 RULES

```

**Advanced
Technologies
Division**

3.3 INQUIRE

The INQUIRE function permits the user to examine the knowledge base in any desirable level of detail. The user can examine the decision structure (the informatin entered during the BUILD function), examine any or all rules, or obtain a summary of the rules in the knowledge base.

GRC

```
*MAIN MENU* :   B BUILD      T TRAIN      E EXERCISE
TYPE LETTER :   I INQUIRE   G GENERALIZE  C COMPRESS
+ FILE NAME :   M MODIFY     N CHECK CONSISTENCY P CHECK COMPLETENESS
AND RETURN :    R RECOVER    S STOP
```

SELECTION? P ←

THE FOLLOWING CASE IS A DISTANCE OF 0.452 FROM THE NEAREST ENTRY
IN THE KNOWLEDGE BASE:

```
PART                = B
REPAIRABLE           = YES
PRIORITY              = 4
REPAIR COST           = 782
CUMULATIVE REPAIR COST = 1561
TIME IN SERVICE       = 12
```

DO YOU WANT TO ADD THIS CASE AS A NEW TRAINING CASE (Y OR N)? Y ←

REPAIR OR REPLACE = \$BSTOCK ←

DO YOU WANT TO CONTINUE (Y OR N)? Y ←

Advanced
Technologies
Division

3.4 CHECK COMPLETENESS

This function allows the expert to identify areas in the domain for which TIMM™ has no rules that could be applied. Using these attributes of the decision factors, TIMM™ identifies specific situations that are "most dissimilar" from the rules in the knowledge base. These are presented as training examples for the expert.



MAIN MENU : B BUILD T TRAIN E EXERCISE
TYPE LETTER : I INQUIRE G GENERALIZE C COMPRESS
+ FILE NAME : M MODIFY N CHECK CONSISTENCY P CHECK COMPLETENESS
AND RETURN : R RECOVER S STOP

SELECTION? G ←

I WILL NOW ATTEMPT TO FORMULATE GENERAL RULES, BASED ON THE KNOWLEDGE IN THE
CURRENT KNOWLEDGE BASE.

DO YOU WANT ME TO PRESENT THE POTENTIAL NEW RULES TO YOU FOR APPROVAL OR
DISAPPROVAL (Y OR N)? Y ←

OK, LET'S BEGIN. TO STOP FOR ANY REASON, ENTER "\$".

CONSIDER THE FOLLOWING RULE:

IF:

PART = *
REPAIRABLE = *
PRIORITY = *
REPAIR COST = >101
CUMULATIVE REPAIR COST = >660
TIME IN SERVICE = >74

THEN:

REPAIR OR REPLACE = \$B\$TOCK(100)

IS THIS RULE CORRECT (Y, N, OR ?)? N ←

Advanced
Technologies
Division

3.5 GENERALIZE

The GENERALIZE function searches the existing knowledge base and attempts to formulate general rules from the more specific examples and rules in the knowledge base. This is one of the more powerful functions available in TIMM™ making it possible to induce general rules from specific examples.

The rules are presented to the expert for acceptance or rejection. If the expert is not sure of a rule (?), TIMM™ will not add the rule to the knowledge base, but searches for similar rules that might be more acceptable.



```
*MAIN MENU* : B BUILD      T TRAIN      E EXERCISE
TYPE LETTER : I INQUIRE   G GENERALIZE  C COMPRESS
+ FILE NAME : M MODIFY     N CHECK CONSISTENCY P CHECK COMPLETENESS
AND RETURN. : R RECOVER    S STOP
```

SELECTION? C ←

MY ATTEMPT TO COMPRESS THE KNOWLEDGE BASE IS COMPLETE.

**Advanced
Technologies
Division**

3.6 COMPRESS

The COMPRESS function is used after GENERALIZing. It deletes specific examples that are covered by more general rules. This keeps the knowledge base as small as possible (for efficiency) without losing information. The expert can use INQUIRE in conjunction with COMPRESS to see which rules were deleted from the knowledge base.



MAIN MENU : B BUILD T TRAIN E EXERCISE
TYPE LETTER : I INQUIRE G GENERALIZE C COMPRESS
+ FILE NAME : M MODIFY N CHECK CONSISTENCY P CHECK COMPLETENESS
AND RETURN : R RECOVER S STOP

SELECTION? N ←

I WILL NOW LIST FOR YOU ALL RULES AND PAIRS OF RULES IN MY KNOWLEDGE BASE WHICH CAN LEAD TO INCONSISTENT DECISIONS. IF YOU HAVE NOT ALREADY DONE SO, I SUGGEST THAT YOU PRINT THE KNOWLEDGE BASE AFTER I AM FINISHED SO YOU CAN REFER TO IT.

SHALL I COPY THE INFORMATION INTO A FILE IN YOUR DIRECTORY AS I DISPLAY IT (Y OR N)? N ←

INCONSISTENCIES IN THE KNOWLEDGE BASE FOR THE EXPERT SYSTEM ORGANIZITIN MAINTENANCE"

THE KNOWLEDGE BASE APPEARS TO CONTAIN NO INCONSISTENCIES.

Advanced
Technologies
Division

3.7 CHECK CONSISTENCY

The CHECK CONSISTENCY function examines the knowledge base for rules that would result in contradicting decisions in a specific case. Rules that have more than one choice or rules with overlapping IF clauses and distinct THEN clauses are identified. The expert has the option of deleting rules or preserving the "inconsistency" if it is appropriate.

Inconsistency often arises when multiple experts train an expert system. This function identifies where the experts disagree. In some cases, reasonable men may reasonably disagree, and the expert system can make a decision that includes multiple choices.

GRC

```
*MAIN MENU* : B BUILD T TRAIN E EXERCISE
TYPE LETTER : I INQUIRE G GENERALIZE C COMPRESS
+ FILE NAME : M MODIFY N CHECK CONSISTENCY P CHECK COMPLETENESS
AND RETURN : R RECOVER S STOP
```

SELECTION? M ←

```
*MODIFY* : AF ADD A FACTOR AC ADD A CHOICE R RETURN TO
TYPE LETTERS : DF DELETE A FACTOR DC DELETE A CHOICE MAIN MENU
AND RETURN DR DELETE A RULE N CHANGE A NAME
```

SELECTION? DR ←

TELL ME THE RULE YOU WANT TO DELETE: 65 ←

RULE 65

IF:

```
PART = C
REPAIRABLE = YES
PRIORITY = 2
REPAIR COST = 404
CUMULATIVE REPAIR COST = 1459
TIME IN SERVICE = 181
```

THEN:

```
REPAIR OR REPLACE = $BSTOCK(100)
```

IS THIS THE RULE YOU MEAN (Y OR N)? N ←

Advanced
Technologies
Division

3.8 MODIFY

The MODIFY function allows the user to make additions, deletions, and changes to the decision structure and to the knowledge base. Decision choices and factors can be added, deleted, or given new names. Rules can be deleted. Spellings and prompts can be changed.


```

*MAIN MENU*   :   B BUILD       T TRAIN           E EXERCISE
TYPE LETTER   :   I INQUIRE    G GENERALIZE      C COMPRESS
+ FILE NAME   :   M MODIFY      N CHECK CONSISTENCY P CHECK COMPLETENESS
AND RETURN*   :   R RECOVER     S STOP

```

SELECTION? E ←

TO EXERCISE THE EXPERT SYSTEM, I NEED VALUES FOR THE FACTORS YOU'VE DEFINED. EITHER YOU PICK THEM, OR I WILL. IF YOU PICK THEM, YOU MAY GIVE A SINGLE VALUE FOR A FACTOR, OR AS MANY AS TEN VALUES WITH CERTAINTY NUMBERS (WHOLE NUMBERS UP TO 100), LIKE THIS:

```

      TEMPERATURE = HIGH   (50)
                  NORMAL   (50)

```

/ (TO INDICATE THAT'S ALL)

YOU CAN ALSO GIVE THE VALUE OF A FACTOR AS "*" FOR "IRRELEVANT" OR "?" FOR "DON'T KNOW". RANGES OF VALUES AREN'T ALLOWED. TO STOP, ENTER A VALUE AS "\$".

IF YOU ASK ME TO NARROW DOWN THE CHOICES AS I GO, I'LL LIST MY CHOICES AFTER EACH FACTOR VALUE IS SPECIFIED, REVISING MY LIST WITH EACH NEW FACTOR VALUE. OTHERWISE, I'LL WAIT FOR ALL THE FACTOR VALUES AND THEN MAKE A DECISION. I'LL SHOW MY CHOICES WITH CERTAINTY NUMBERS, AND ALSO A RELIABILITY NUMBER TO LET YOU KNOW (ON A SCALE OF 100) HOW CONFIDENT I AM IN MY DECISION.

AFTER EACH TEST CASE, I'LL ASK YOU IF MY DECISION WAS CORRECT (Y, N, OR E). "E" MEANS EXPLAIN -- I'LL SHOW YOU WHICH RULES I USED TO MAKE MY DECISION. THEN I'LL ASK IF YOU WANT ANOTHER TEST CASE (Y, N, OR A). "A" MEANS ADD THE PREVIOUS TEST CASE TO THE KNOWLEDGE BASE AS A NEW RULE. I'LL GIVE YOU THE ADD OPTION ONLY IF ALL FACTORS WERE GIVEN DEFINITE VALUES.

**Advanced
Technologies
Division**

3.9 EXERCISE

The EXERCISE function allows a non-expert user to access the knowledge base, ask questions, and receive answers. Several options are available.



TIMM™ "EXERCISE" FUNCTION (CONT.)

SHALL I PICK THE TEST CASES FOR YOU (Y OR N)? Y ←

SHALL I NARROW DOWN THE POSSIBLE CHOICES AS I GO (Y OR N)? N ←

DO YOU WANT ME TO USE TERSE (T) OR VERBOSE (V) MODE? T ←

FOR THIS CASE, SHALL I EXCLUDE ANY CHOICES (Y OR N)? N ←

PART	= A
REPAIRABLE	= YES
PRIORITY	= 1
REPAIR COST	= 755
CUMULATIVE REPAIR COST	= 1104
TIME IN SERVICE	= 24
REPAIR OR REPLACE	= \$BSTOCK(100)
	(RELIABILITY = 100)

DOES THIS CONCLUSION SEEM REASONABLE (Y, N, H, OR E)?

Advanced
Technologies
Division

The first example shows a case in which TIMM™ has generated a situation at random and given a decision. TIMM™'s reliability factor tells, on a scale of 0 to 100, how well the example matches the rule employed to make the decision. In this case, the test case exactly matches a rule.

Several options are available. The user may ask for an explanation (E) in which case the rules used to make the decision will be displayed. If the user answers No to the question, TIMM™ can add the case to the knowledge base and prompt the user for the correct decision. Entering (H(elp)) will cause TIMM™ to display the help prompt for the decision.



SHALL I PICK THE TEST CASES FOR YOU (Y OR N)? N ←

SHALL I NARROW DOWN THE POSSIBLE CHOICES AS I GO (Y OR N)? N ←

DO YOU WANT ME TO USE TERSE (T) OR VERBOSE (V) MODE? V ←

FOR THIS CASE, SHALL I EXCLUDE ANY CHOICES (Y OR N)? N ←

FOR THIS CASE, DO YOU WANT TO SPECIFY ANY FACTOR VALUES FIRST (Y OR N)? N ←

WHICH PART FAILED?

ANSWER: A ←

IS THE FAILED PART REPAIRABLE?

ANSWER: YES ←

WHAT IS THE PRIORITY OF THE FAILED PART?

ANSWER: 3 ←

WHAT IS THE ESTIMATED COST TO REPAIR THE FAILED SUBASSEMBLY?

ANSWER: 585 ←

HOW MUCH HAS BEEN SPENT IN PRIOR REPAIRS FOR THE FAILED PART?

ANSWER: 585 ←

HOW LONG HAS THE FAILED PART BEEN IN SERVICE?

ANSWER: 140 ←

REPAIR OR REPLACE:

DETERMINE IF A REPLACEMENT FOR THE FAILED PART IS IN BENCH STOCK.

(100)

RELIABILITY = 77

**Advanced
Technologies
Division**

If the user selects the verbose mode of operation, the verbose prompts are used by TIMM™ to solicit factor values. In this case, the specific situation is not exactly like the closest rule, so TIMM™ reports a reliability less than 100.

4. SUMMARY

TIMM™ is a flexible and practical system for developing and using expert systems. TIMM™ allows experts to interact directly with the program in constructing knowledge bases by automating the task of knowledge acquisition. TIMM™ interviews the expert to define the domain, poses situations for the expert to consider, induces general rules from specific cases, and can include expertise from multiple experts.

Rapid prototyping and iterative development of expert systems are supported by TIMM™. The basic TIMM™ functions can be viewed as a toolbox and each of the tools can be used to perform a task during system development. The ability to modify and refine decision domains and rules allows basic systems to be constructed quickly. By linking multiple expert systems together, both structure and hierarchy can be improved on the overall system design.

TIMM™ is economical. TIMM™ is written in Fortran and runs in standard computing environments ranging from IBM PC class computers to large mainframes. No special programming is required to input expert systems using TIMM™. Thus, the critical area of knowledge acquisition can be addressed directly.

ACKNOWLEDGMENTS

The author is indebted to many individuals at GRC who have contributed to the development of TIMM: Mr. Dennis Cooper, Dr. Lester Ford, Dr. Stephen Kiselewich, Mr. Jim Kornell, Dr. Richard Parker, Dr. Ralph Pennington, Mr. Phil Sewell, Mr. Randall Urban, and Mr. Rollin Weeks. Dr. Kiselewich is the principal author of the TIMM™ program. The Organizational Maintenance expert system described in this paper was developed by Mr. Gerald Paul and Dr. Peter McWhite.

REFERENCES

1. Feigenbaum, E. A., "Knowledge Engineering for the 1980's," Computer Science Department, Stanford University, 1982.
2. Schank, R., "The Current State of AI: One Man's Opinion," AI Magazine, Vol. 4, No. 1, Winter/Spring 1983.
3. Norman, D., "Twelve Issues for Cognitive Science," Cognitive Science, Vol. 4, No. 1, January-March 1980.
4. Feigenbaum, E. A., P. McCorduck, Fifth Generation, Addison-Wesley, Reading, Mass., 1983, pp. 80-84.
5. Davis, R., "Expert Systems: Where Are We?" And Where Do We Go From Here?," AI Magazine, Vol. 3, No. 2, Spring 1982.
6. Hayes-Roth, F, D. Waterman, and D. Lenat, Building Expert Systems, Addison-Wesley, Reading, Mass., 1983.
7. Wertheimer, M., Productive Thinking, Harper Torchbooks, New York, NY, 1959.
8. Lakoff, G. and M. Johnson, "The Metaphorical Structure of the Human Conceptual System," Cognitive Science, Vol. 4, No. 2, April-June 1980.
9. Gentner, D., "Structure-Mapping: A Theoretical Framework for Analogy," Cognitive Science, Vol. 7, No. 2, April-June 1983.
10. Wittgenstein, L., Philosophical Investigations, trans. G.E.M. Anscombe, Blackwell, Oxford, England, 1953.
11. Aikens, Jr., "Prototypical Knowledge for Expert Systems," Journal of Artificial Intelligence, Vol. 20, No. 2, Feb. 1983.
12. Friedland, P., "Acquisition of Procedural Knowledge from Domain Experts," Computer Science Department, Stanford University, Stanford, CA 94305.
13. Minsky, M., "A Framework for Representing Knowledge," in The Psychology of Computer Vision, P. Winston, Ed., McGraw-Hill, New York, NY, 1975.
14. Schank, R., R. Abelson, Scripts Plans Goals and Understanding, Lawrence Erlbaum Associates, Hillsdale, NJ, 1977.
15. Winograd, T., "Frame Representations and the Declarative/Procedural Controversy," in Representation and Understanding, D. Bobrow, A. Collins, eds., Academic Press, New York, NY, 1975.

16. Clancy, W., review of States of Mind by Jonathon Miller, AI Magazine, Vol. 4, No. 4, Winter 1983.
17. Pentland, A. and M. Fischler, "A More Rational View of Logic," AI Magazine, Vol. 4, No. 4, Winter 1983.
18. Hayes-Roth, F., "The Role of Partial and Best Matches in Knowledge Systems," in Pattern-Directed Inference Systems, D. Waterman, F. Hayes-Roth, Eds., Academic Press, New York, NY, 1978.
19. Kornell, J., "Embedded Knowledge Acquisition to Simplify Expert Systems Development," Proceedings of the 1984 Conference on Intelligent Systems and Machines, Detroit, Mich., 1984.
20. Hendrix, G. G., E. D. Sacerdoti, D. Sagalowicz, and J. Slocum, "Developing a Natural Language Interface to Complex Data," ACM Transactions on Database Systems, Vol. 3, No. 2, June 1978, pp. 105-147.

**DECISION SUPPORT AND EXPERT SYSTEMS
FOR EMERGENCY MANAGEMENT OPERATIONS: A MICROCOMPUTER APPROACH**

By

FRANK DE BALOGH, Ph.D.

Director, Decision Support Systems Laboratory
Institute of Safety and Systems Management
UNIVERSITY OF SOUTHERN CALIFORNIA
Los Angeles, California 90089-0021

1. INTRODUCTION.

Using micro-computer technology to improve the productivity and capacity of various facets of our economy and government infrastructure is now a major thrust in many professions including emergency management. This paper is designed to: (1) provide an overview of some of the activities of the University of Southern California (USC) in using micro-computers in emergency management to develop greater capacity in decision support, and (2) to make recommendations to the Federal Emergency Management Agency (FEMA) as to how it should seek to take advantage of the opportunity offered by a new class of information technology - Expert Systems (ES) - to advance the field of emergency management.

Since the two USC research examples to be described in this paper are incorporated in a particular class of information system - Decision Support Systems (DSS) - this concept will be addressed first. The term "DSS" is today widely employed in financial and business circles and is slowly finding its way into the vocabulary of emergency management. A Decision Support System (DSS) is a collection of computational, analytical and decision-theoretic models together with a data base which reside in an interactive computer environment. Its express objective is to improve the effectiveness of what have been termed knowledge workers. A DSS consists of three major components: (1) the Data Base, (2) the Model Base, and (3) the software system which contains the Data Base Management System, the Model Base Management System and the Dialog Management System. The computer hardware and software, which make up these components, (as a system) is called the DSS Generator (DSSG) since it is the vehicle used to develop (generate) specific applications in emergency operations (Sprague and Carlson, 1982).

To provide a clear understanding of the focus of USC's research program, which also underlies this paper the following summary is provided. The majority of USC's research effort is based on the premise: CAPACITY BUILDING THROUGH WIDESPREAD TECHNOLOGY TRANSFER. That is why we are employing widely available commercial hardware and software such as the IBM PC AT, LOTUS' 1-2-3, Microrim's CLOUT, etc., in applications development. FEMA has stated that it is turning to the off-the-shelf commercial sector to meet some of its hardware and software needs (Giuffrida, 1985; Campbell, 1984).

Organizationally, USC has the following components dealing with emergency management:

*** DECISION SUPPORT SYSTEMS LABORATORY (DSSL):**

The function of this laboratory, one of the few such university facilities in the United States, is to conduct state-of-the-art research as well as to develop prototypes in the field of decision support systems (DSS). This research effort is designed to develop DSS applications and programs for a wide range of sponsors and users such as government agencies, aerospace companies, business, and industrial organizations. Such DSS can help managers make better decisions in complex, unstructured, or semi-structured situations. Current research projects focus on the development of (1) an Earthquake Decision Support System to support earthquake management efforts and projects designed to study the impact of hardware and software configurations and presentations in decisionmaking capabilities; (2) a decision support system to aid in the management of hazardous materials spills or discharges, and, (3) computerized management aids for emergency evacuations of smaller cities or localities. The DSSL is currently working on non-disaster related systems as well.

*** RISK AND EMERGENCY MANAGEMENT LABORATORY (REML):**

Work to date by this laboratory has included projects done in (dam safety) risk assessment and addressing the needs of special populations during a disaster, etc. REML and others at USC have also been involved in a project to assist public officials (specifically, EPA policymakers) on how to communicate risks to public officials and the public at large.

Also noteworthy is the SPECIAL ISSUE ON EMERGENCY MANAGEMENT in PUBLIC ADMINISTRATION REVIEW (PAR) published in January 1985 funded by FEMA, edited by William Petak, REML director. The article on decision support systems in that special issue, written by William Wallace of Rensselaer Polytechnic Institute and this author, reflect the same premise that capacity building requires ubiquitous hardware and software (Petak, 1985; Wallace and de Balogh, 1985).

*** COMEX PROJECT:**

This is a nationally-recognized computer simulation and training laboratory for public administrators organized within the USC School of Public Administration with support from IBM. It teaches the dynamics of policymaking and planning by using a model of a medium-sized city which is driven by a data base on a minicomputer. Prospective policies ("What Ifs") are tried out through role playing at terminals using various scenarios.

The next section will focus on USC research efforts underway in two areas: (1) EARTHQUAKE DECISION SUPPORT SYSTEM (EDSS) - a two year project funded by the National Science Foundation using two main-frame software packages (ADMATCH 80 and DATAPLOT). This is being carried out in conjunction with USC's Social Science Research Institute (SSRI) and DSSL, and (2) HAZARDOUS MATERIALS (TOXDSS) - a DSSL project using the LOTUS 1-2-3 software and microcomputers.

2. DECISION SUPPORT SYSTEMS ARCHITECTURE FOR EARTHQUAKE MANAGEMENT.

2.1 INTRODUCTION.

The EDSS prototype is designed to extend the capability of planners and public officials to make effective decisions in mitigation and other stages of the earthquake life cycle. EDSS is a combination of data base files, decision aiding models, and user communications which facilitate employment by a variety of users (See Figure 1).

It consists of several basic files on building characteristics, population distribution, quake intensity, and upgrade costs. These are used by various mathematical models that generate damage assessments as well as the costs of upgrading building stock by regional (planning) areas which divide a locality. The user employs a series of help menus as well as the ad hoc reporting capabilities of the system (Schuerman, Petak and de Balogh, 1982).

2.2 GENERAL EDSS FUNCTIONAL GOALS.

EDSS is capable of aiding in the following mitigation planning functions:

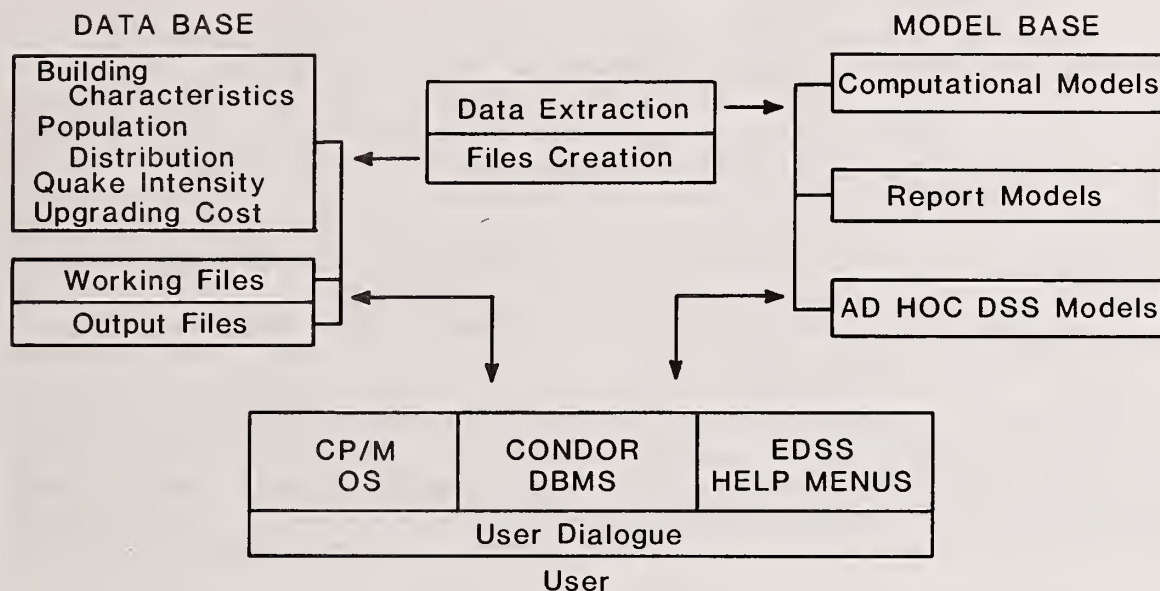
*** ASSESS THE IMPACT OF AN EARTHQUAKE ON:**

- + Present buildings in a planning area
- + Proposed buildings
- + Buildings with proposed upgrading

*** ESTIMATE THE COST OF:**

- + Proposed improvements in present buildings to 1934 standards (in the Los Angeles city prototype)

Figure 1
EDSS CONCEPTUAL FRAMEWORK



Sample EDSS Display Output of Simulated Earthquake Damage by Planning Area
(City of Los Angeles)

Quake Damage Assessment Using the Current Building Stock

Central Business District Quake Magnitude = 6.3 Date:

Results for Structure Class 3 and Age 3

Square Feet Damaged and Replacement Cost in Each Planning Area

Planning Area	Sq. Ft. Damaged (in thousands)	Replacement* Cost (in millions)
PLNAREA NAME	VALL1 (Subtotal)	VALLST (Subtotal)
103 SOUTHEAST LA	0	0
104 W.ADAMS, LEIMERT, BALDWIN	25	21
105 S.CENTRAL LA	0	1
106 WILSHIRE	58	107
107 HOLLYWOOD	35	12
109 WESTLAKE	2	0
110 CENTRAL CITY	100	192
326 WESTWOOD	45	14
327 W.LA, CENTURY-CITY, R.PK	21	38
328 PALMS, MAR-VISTA, DEL-REY	4	3
330 WESTCHESTER, PLAYA-D-REY	8	3
Total	298	391

*Based upon assessed value.

- + To 1974 standards (Los Angeles prototype)
- * ESTIMATE THE POTENTIAL SAVINGS FROM IMPROVEMENTS TO:
 - + Present buildings
 - + Future Buildings
- * PROVIDE SUMMARY STATISTICS ON:
 - + Total damage
 - + Impacted geographic areas
 - + Building categories
- * CONDUCT 'WHAT IF' ANALYSIS ON THE ABOVE BY VARYING:
 - + Cost per square foot to upgrade buildings
 - + Earthquake intensity

For policy simulation purposes, the decisionmaker provides ground shaking intensity data for each planning area which is used by EDSS to calculate percentage of square feet damage and percentage of assessed value lost for each building classification and age group (de Balogh, Petak and Sessler, 1985).

2.3 FUNCTIONS OF SELECTED EDSS MODULES.

The system contains two major modeling components:

* DAMAGE ASSESSMENT - this permits the calculation of damages to certain building categories within a planning area given certain earthquake intensities which are input into the system by the user.

* BUILDINGS UPGRADING - this calculates the cost of upgrading certain age classified buildings within a planning area from one standard to another. The cost per square foot of upgrading a particular category is entered prior to any calculations by the user.

Both models permit "before" and "after" building upgrade damage assessment calculations and loss value estimation. Several of the key externally-provided data items that the EDSS user may define and vary are:

- > COST PER SQUARE FOOT TO UPGRADE BUILDINGS -
this can range from \$2.00 to \$15.00 depending on building category;
- > SPATIAL INTENSITY OF THE EARTHQUAKE -
this can range from 5.0 through 9.0 on the Richter scale.

The model calculates percent loss according to a Modified Mercalli Index (EDSS takes this from Steinbrugge but without the "knee" shown in his vectors) (Algermissen and Steinbrugge, 1980).

3. DECISION SUPPORT SYSTEMS ARCHITECTURE FOR HAZARDOUS MATERIALS MANAGEMENT.

3.1 INTRODUCTION.

The following describes a prototype Hazardous Materials DSS (TOXDSS) which focuses both on preparedness (through training support) and actual response to a chemical spill of limited size either in a city/county or on a military base while the substance was in transport. The purpose of TOXDSS is to aid managers in configuring the best composition of a response team in terms of specialist staffing, vehicles, equipment, protective clothing, clean-up and neutralization chemicals and other requirements. It consists of (1) a data base of emergency personnel and their qualifications, the types of toxic chemicals and where they are stored, the characteristics of these chemicals, and, the types of requirements associated with spill management, (2) a model base which includes predetermined mixes of personnel, equipment, etc., needed to field a response team for a given type and magnitude of spill, and (3) a menu-driven system for developing a response team based upon input provided by the user as to the type, size, location, etc., of the spill.

This DSS is a demonstration model first developed at USC (Germany campus) in Stuttgart, West Germany by a team including engineers and programmers for potential application in U.S. Army, Europe (USAREUR), and later expanded in Los Angeles.

3.2 GENERAL TOXDSS FUNCTIONAL GOALS.

The prototype incorporates a number of unique features as well as some of the functions that are planned for future modules. (Hearne, et al., 1984); (See Figure 2).

These functions are:

- + Determine the composition of the incident response team to include in-house as well as contract resources;

- + Assist in providing immediate health and safety information to those reporting a spill as necessary;

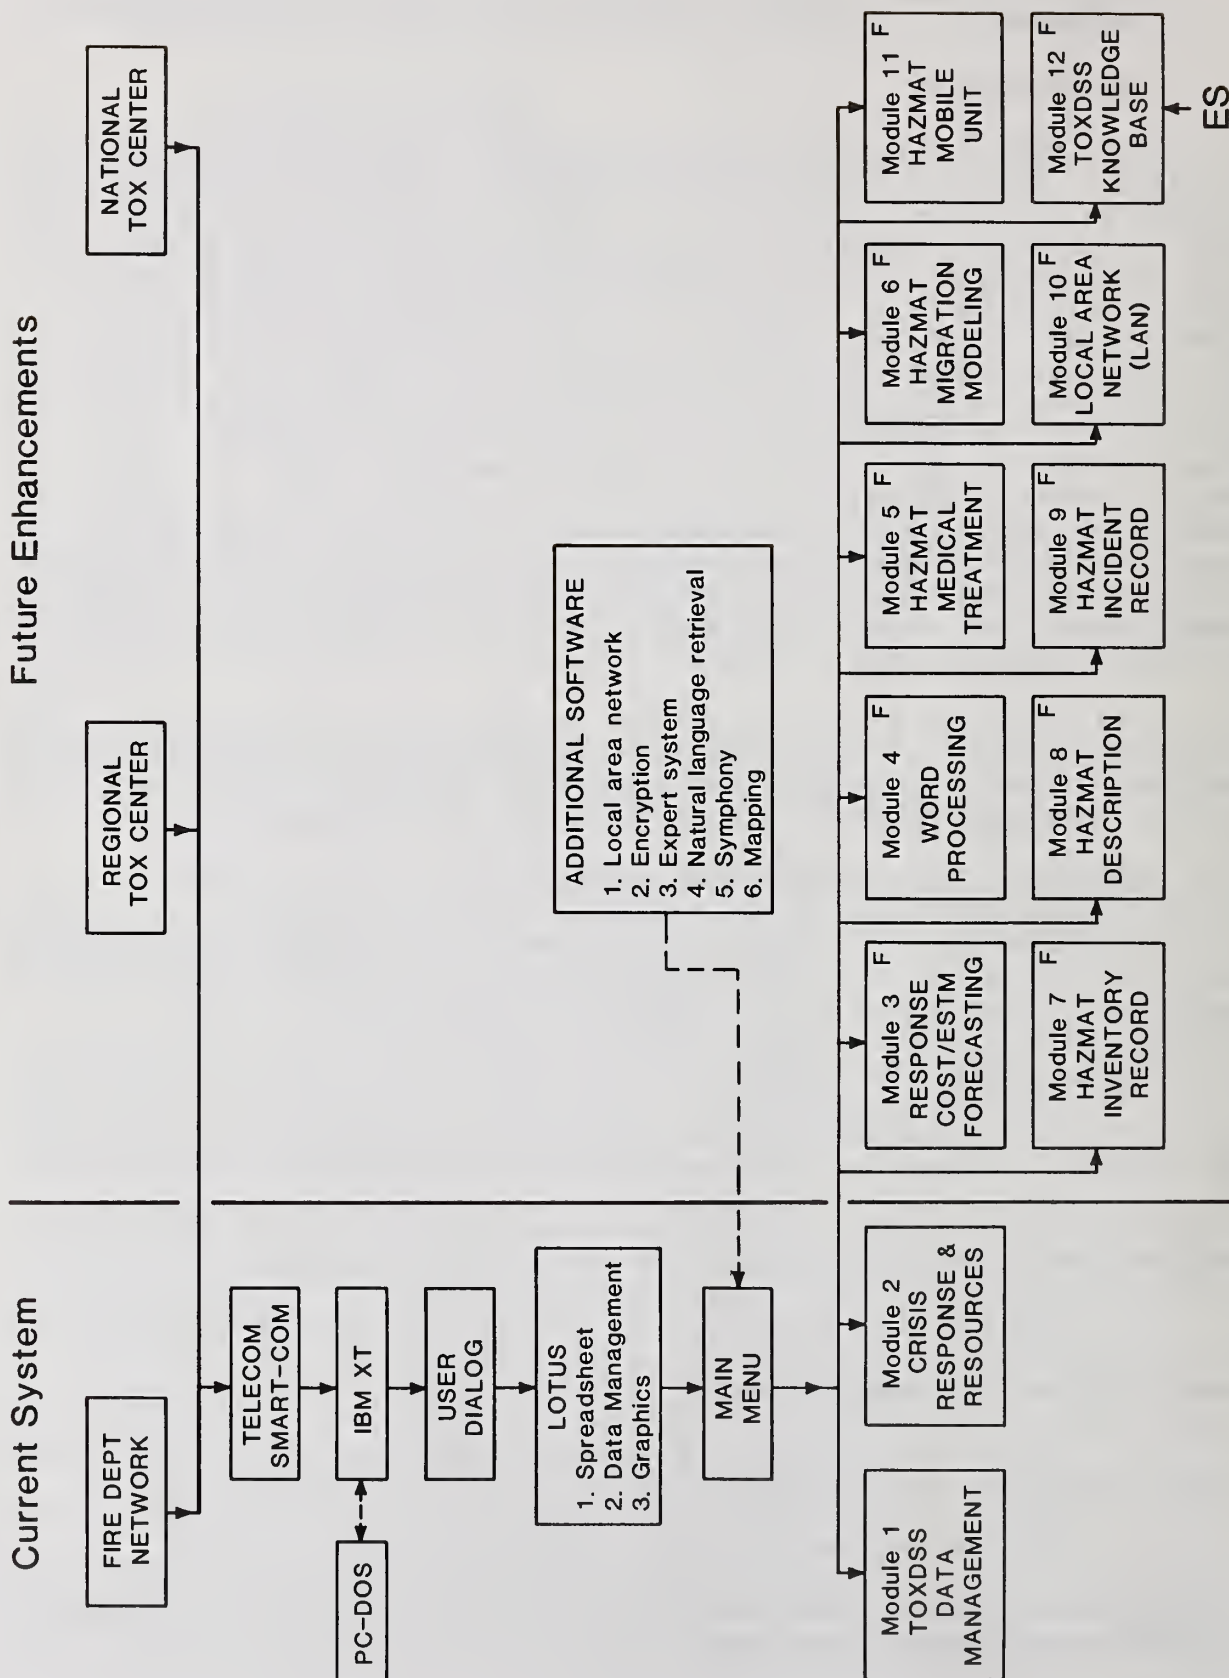
- + Notify appropriate agencies about the spill, its environmental impact, potential for migration, and other consequences and requirements;

- + Identify possible disposal sites given the specific characteristics of the spill;

- + Aid in determining spill responsibility;

- + Estimate immediate spill response costs for budgetary planning and cost recovery.

Figure 2
TOXDSS SYSTEM FRAMEWORK MODULES



3.3 FUNCTIONS OF SELECTED TOXDSS MODULES.

The three TOX DSS components - data base, models, and user dialog - for this module are discussed in more detail. However, since it is a prototype the hazard migration models referred to are not fully developed (de Balogh, 1985).

TOXDSS incorporates a model similar to that found in the National Oil and Hazardous Substances Contingency Plan (NCP). (See Environmental Protection Agency, CFR, in references.) It employs a system for ranking hazardous materials based on a mathematical model with simulation capabilities that uses the following variables:

- + CHEMICAL CHARACTERISTICS VALUES (CHEMCHAR) - the weighted values are taken directly from the NCP and include ignitability, reactivity, and toxicity/persistence;

- + INCOMPATIBILITY SCORES (INVNTORY) - weighted values taken from the NCP;

- + ACCESSIBILITY SCORES (INVNTORY) - weighted values taken from the NCP;

- + POTENTIAL MIGRATION SCORES (INVNTORY) - given the complexity of modeling this phenomenon, this contains only limited estimates based on a 20 point scoring system (American Petroleum Institute, 1972).

In determining the most cost-effective combination of response resources for a specific incident the module calculates hazard scores for the spill for: (1) fire and explosion; (2) direct human contact; (3) potential migration for ground water, surface water and air. Once the hazard associated with a specific incident has been established the CRTEAM data base is queried to select recommended response staff, equipment, vehicles, etc. The selection is based on a calculated hazard score that is less than the values listed in the CRTEAM data base for the particular item of response equipment or specialist. Size of the spill is also taken into consideration in determining the composition of the incident response team.

The prototype is geared to be used by non-specialist personnel with a minimum of training. This is especially important during the response phase of an incident. Therefore, it has been designed to be menu driven. It prompts the user to enter the chemical's National Stock Number (NSN), quantity spilled, and location. After accessing a number of data bases and assessing hazard scores, the module automatically presents the suggested response team composition. Actual field testing of this module has been initiated by the County of San Diego, California, with the intention of DSSL improving the prototype based upon results from that test.

TOXDSS in its broadest design concept will incorporate a number of additional functional capabilities including a knowledge base (ES) component. A cost estimation and forecasting capability will be developed in the near future. This interfaces with the previous module and takes the per hour cost of all response resources involved in a specific incident plus other cost categories to develop the cost of a spill. This can be used later to recover costs from parties responsible for the spill. An aggregative cost model permits "what if" analyses by varying cost parameters. It sensitizes all those involved in spill management as to the immediate dollar costs of specific incidents and resources used in response.

Modeling the hazard is to be accomplished by a component that will estimate the migration of a hazardous material spill or discharge in terms of its migration potential on land, water or air environments in a specific incident. This is likely to be the most difficult to develop since it is dependent on detailed data for a specific geographic area which in many cases is hard to obtain before an incident.

4. THE ROLE OF EXPERT SYSTEMS IN DSS ARCHITECTURE.

From an architectural point of view there are three basic forms of relationships between ES and DSS in the foregoing disaster situations and most others:

4.1 ES AS A SEPARATE COMPONENT FOR EACH DSS COMPONENT.

Under this concept each DSS component - data base, model base and dialog - would have a separate expert system component incorporated into its design or externally interfaced for quick call-up (See Figure 3).

4.2 ES AS A FOURTH MAJOR COMPONENT IN DSS.

This would be a single, integrated and highly complex component that would interface with the other major components of DSS. It could be accessed from each of the three DSS components (See Figure 4).

The two primary functions of ES in such DSS would be: (1) improve the operation of the DSS and make it easier to use, and (2) bring expertise to bear on the emergency management life cycle.

4.3 REAL-TIME ES SYSTEM WITH ARTIFICIAL INSTINCTS.

Specialized forms of expert system architectures for narrow but highly important problems such as early warning and early response to catastrophic natural or manmade disaster are also

Figure 3

DSS ARCHITECTURE WITH MULTIPLE ES MODULES

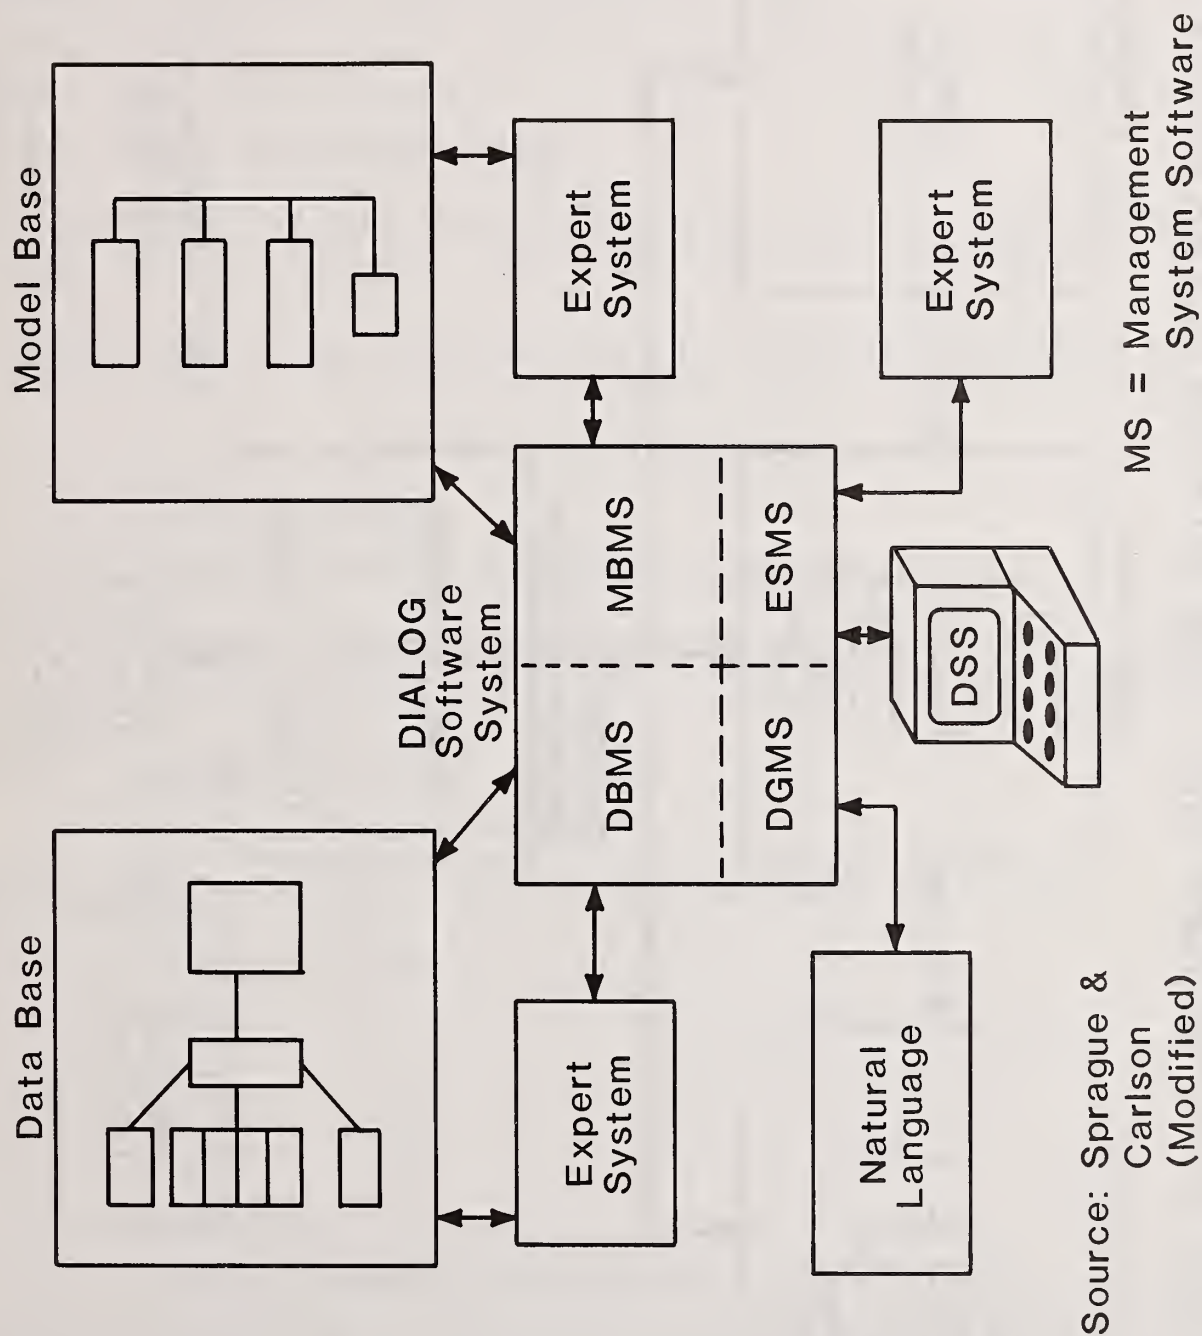
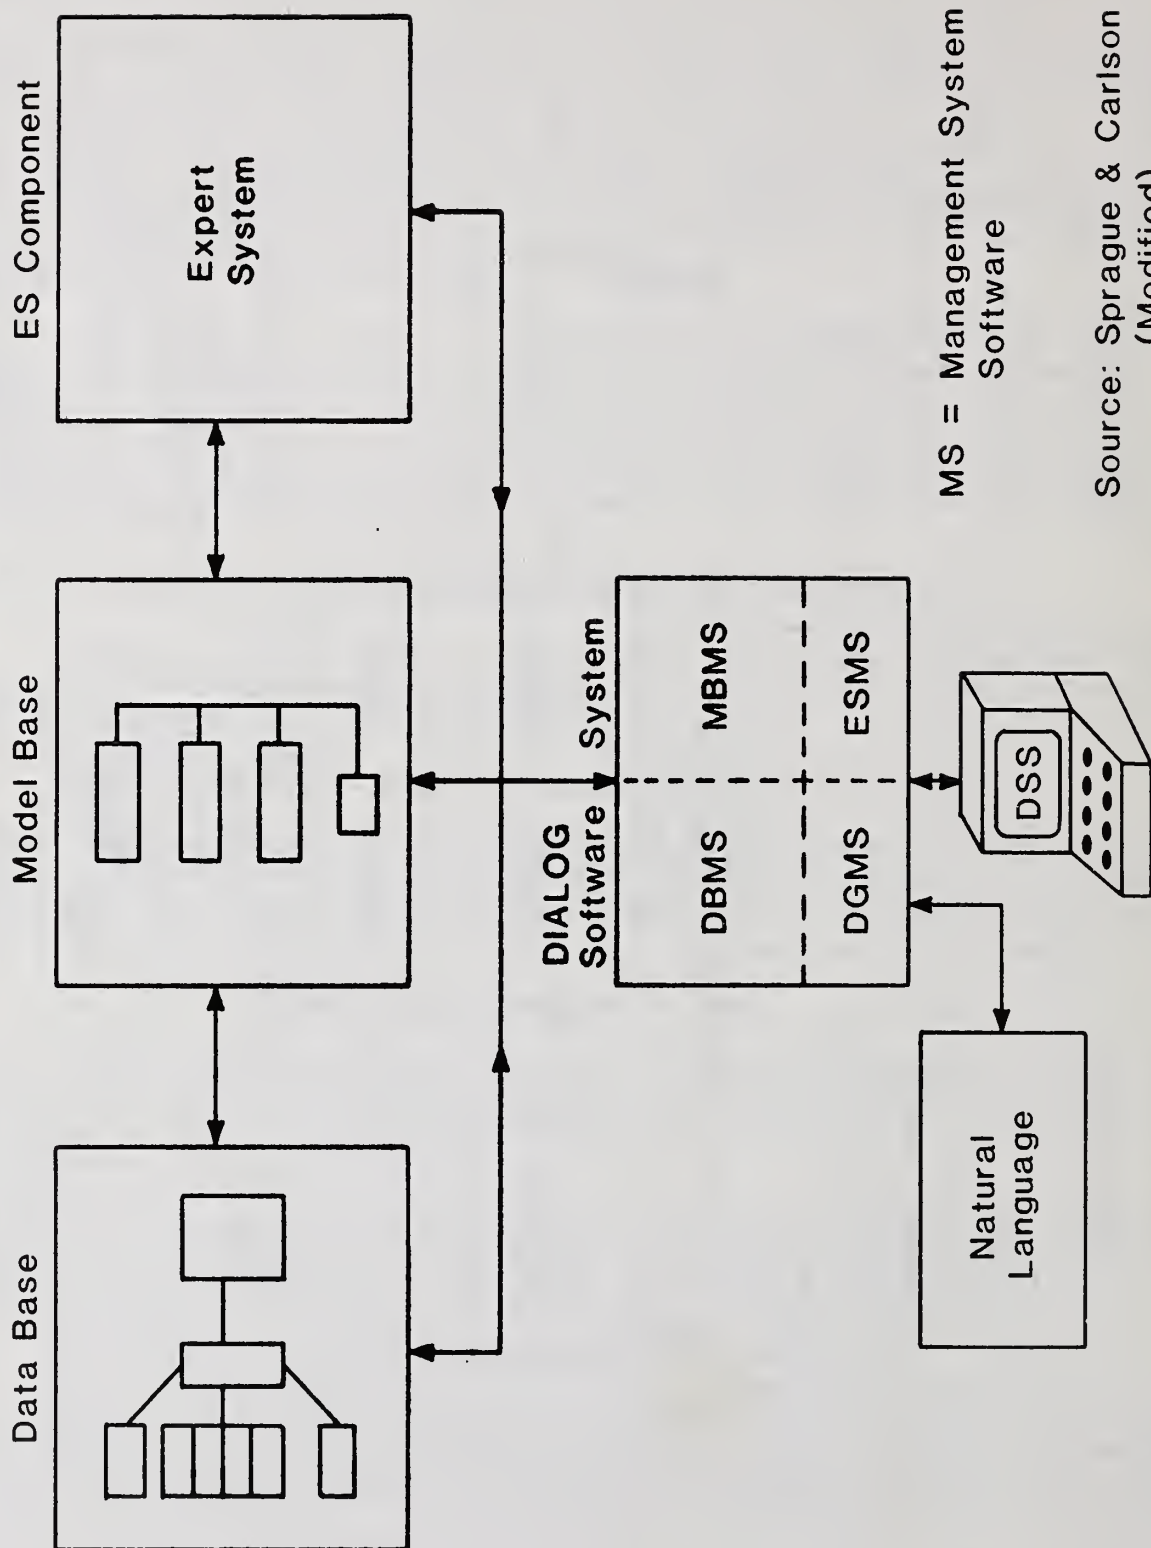


Figure 4

DSS ARCHITECTURE WITH ONE INTEGRATED ES MODULE



forthcoming. These forms are designed for the highly time-dependent, real-time critical applications that are likely to be sensor based. The study of these architectures is drawn from military research and development (R & D) and is termed the (AI)2 concept. This stands for "artificial intelligence - artificial instinct" whose application to emergency management is described in a later section (Gaertner, 1985).

The following section provides more detail about the nature of the DSSG architecture:

5. DSS DIALOG MANAGEMENT SYSTEM.

This is of primary importance since to the user this is "the system." Namely, what the user perceives he is interacting with and how he controls it.

Currently, there are a number of alternative hardware and software modes for controlling a computer system which include:

- * MENU - standard and icon (symbols);
- * COMMAND LANGUAGE;
- * INPUT IN CONTEXT OF OUTPUT (Spreadsheet);
- * SYMBOLIC REPRESENTATION (ICON such as the operating environment of the Apple Mac Intosh or Digital Research's GEM);
- * NATURAL LANGUAGE;
- * WINDOWING (Multi-tasking).

Certain hardware capabilities are also desirable such as:

- * MOUSE OR "NUDGER";
- * COLOR;
- * VOICE;
- * TOUCH;
- * JOYSTICK;
- * OPTICAL CHARACTER RECOGNITION (OCR).

The application of artificial intelligence methodology to facilitate the control (dialogue) of information systems in emergency operations by persons who have little or no formal knowledge for using computers and interacting with complex systems should be a high priority. Improving the ability to use natural English and the specialized terms (jargon) associated with this field to formulate even complex inquiries and better

manage (input/update/merge) emergency related data and files are vital. Special emphasis should be placed on making the dialog with a computer system as simple as possible during the emergency response phase when untrained volunteers may, in some cases, be pressed into service (de Balogh, 1985).

Several micro-computer software products (e.g., Microrim's CLOUT, Excalibur's SAVVY, etc.) are commercially available for this purpose which can now accomodate up to 500 specialized words (synonyms). These can form the basis for further research (Bond, 1985).

A combination of natural language and icon-based interactive dialog environment is essential to assure some user friendliness of such systems (Foster, 1985).

6.0 SMART DATA BASE MANAGEMENT FOR DSS.

An ES component to make data base management systems smart to deal with would include the following functions: (1) simplifying the complexity posed by a large number of emergency operations related files; (2) automatically establishing logical relationships between numerous data variables; (3) assist in determining which files are relevant dependent on the particular stage of an emergency; and, (4) smart updating (self-updating) of selected files in the emergency data base through periodic links with outside systems (e.g., county assessor file, building permits file, etc.) via communication links (Ulmen, 1985).

See Figure 5 for an illustration of the complexity associated with an "ideal" DSS for the earthquake problem.

7. SIMULATION MODELING AND EXPERT SYSTEMS.

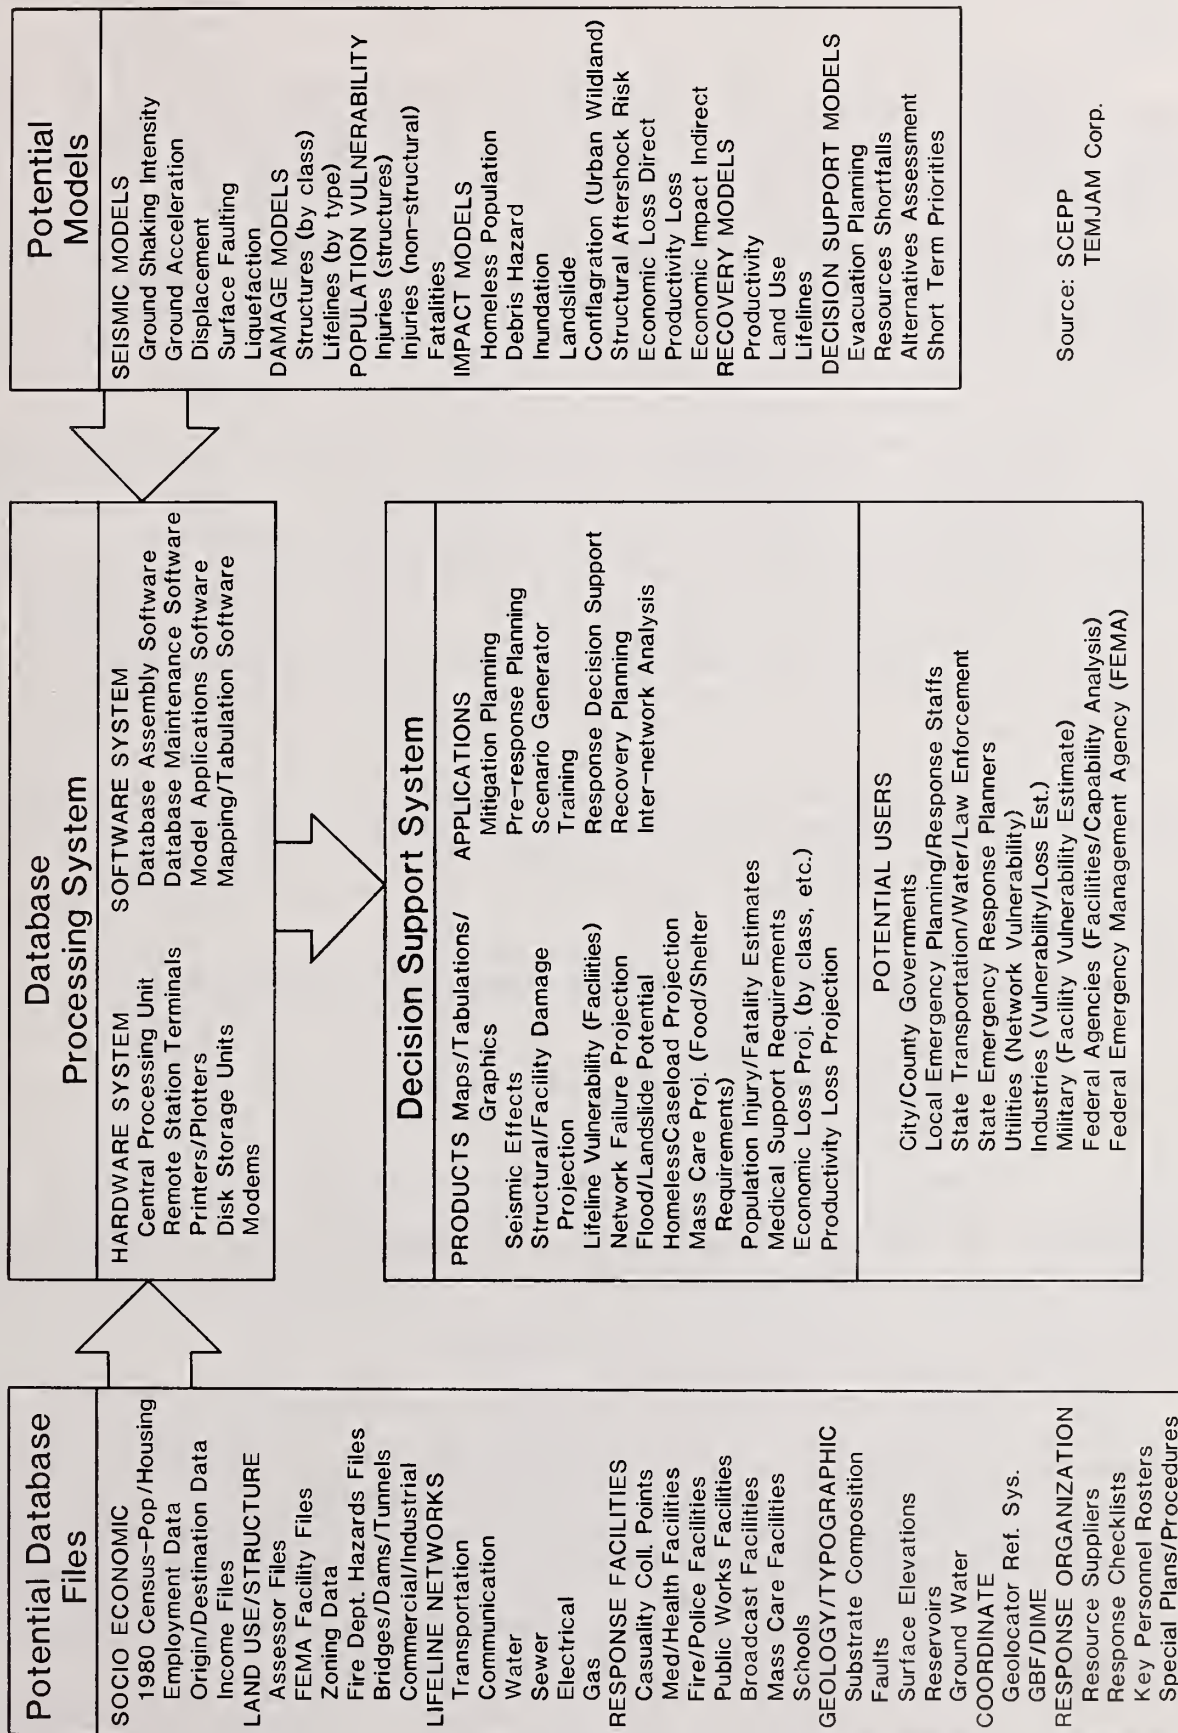
Clearly providing support in the third DSS component for emergency contexts - modeling the hazard and the response - is where ES may have its biggest role in helping to deal with complexity, uncertainty and incomplete methodology.

The following sections discuss the research problems and need for ES within the context of the two previous disaster types.

7.1 EARTHQUAKE IMPACT MODELING.

EDSS illustrates some of the difficulties of modeling in this area especially in terms of the state-of-the-art of current methodology for understanding and forecasting earthquake impact on various categories of buildings. At present, EDSS employs only a simplistic categorization of 5 building types. However, the latest data available on this comes from the work of Applied Technology Council, Inc., which used a Bayesian approach to quantify expert opinion on earthquake impact at different levels on over ninety separate types of buildings (Rojahn, et al, 1984; and Donald Moore,et al., 1985). The accuracy and reliability of

Figure 5
IDEAL DSS CONCEPT FOR EARTHQUAKES



that methodology remains to be demonstrated (Shah and Dong, 1984).

Clearly, as such methodology becomes more complex so does the modeling problem, especially the quality of the data requirements of such models when applied to a particular geographic area such as a medium-to-large size city.

One of the research tasks of the EDSS project is to survey the literature as to the number and types of models found addressing the earthquake problem (Rao, 1985 forthcoming). A finding of the survey is that there are about two dozen such models dealing with seismic risk, mostly on mainframe computers.

They have been developed from many different perspectives: earth sciences, geophysics, geotechnical engineering, structural engineering, urban planning and land use. These separate development efforts show no integration among models. The models are fragmented by discipline and also differ considerably in the level of detail they address. Very few if any can be called user friendly and thus most are not ready for operational use in support of emergency management. At present there is no validated prototype system that is generalizable to public jurisdictions.

Thus, the prospect for capacity building in the earthquake area is not good for the immediate future. It is true that finer estimation methods for seismic risk are emerging and this is where ES may provide opportunities for meaningful advances in the field.

7.2 HAZARDOUS MATERIALS SPILL/DISCHARGE MIGRATION MODELING.

The problem of complexity found in the seismic risk modeling for earthquakes is at least equal if not greater in the hazardous materials area. It is estimated that there are 19,000 separate hazardous materials in existence and a 1000 new materials being added to this number every year. Therefore, the problem of hazard migration modeling is not only horrendous but is dynamic in nature. Any system designed to provide decision support in this area would require a continuous, systematic updating effort of major proportions because it would have to deal with many new substances each year.

The prospect for modeling having a significant impact in the hazardous materials problem is likely to be limited not only for the above reason but also because it is a much newer topic of research interest as compared to the earthquake problem. The most comprehensive ES work relating to the hazardous materials area has been done by the Oak Ridge National Laboratory (Carroll Johnson and Sara Jordan, 1983). That type of capability needs to be made available on a micro-computer before society's capacity to deal with this problem is impacted.

8. THE ARTIFICIAL INTELLIGENCE - ARTIFICIAL INSTINCT (AI)2 CONCEPT.

Recent developments in military R & D, specifically through projects funded by the U.S. Army Signal Warfare Laboratory, have led to the emergence of what may be termed an automated, "real-time" expert system, a specialized Artificial Intelligence architecture. Specifically, this has important implications for automated emergency indicators, warning and response systems that either are or will in the next few years become technologically feasible.

Such systems are designed to be sensor-based to monitor conditions within a complex environment and communicate a warning plus initiate immediate societal response within seconds after such a condition has been identified. They incorporate the expertise of decisionmakers within the system. They rely on a combination of specialized ES software and very fast computer processors linked with reliable telecommunications to a jurisdiction's command and control or emergency operations center and numerous other potentially affected organizations.

Such (AI)2 systems are expected to execute actions within milliseconds (thousandths of a second) by relying on precomputed full or partial response arrays which in turn set into motion a series of automated warning and response mechanisms at various levels of governments involved in the crisis.

The nature of (AI)2 architecture differs somewhat from those previously discussed because: (1) it is much more costly, requiring significant hardware, software, and staffing investments, and (2) contains both real-time and non-real-time components (See Figure 6).

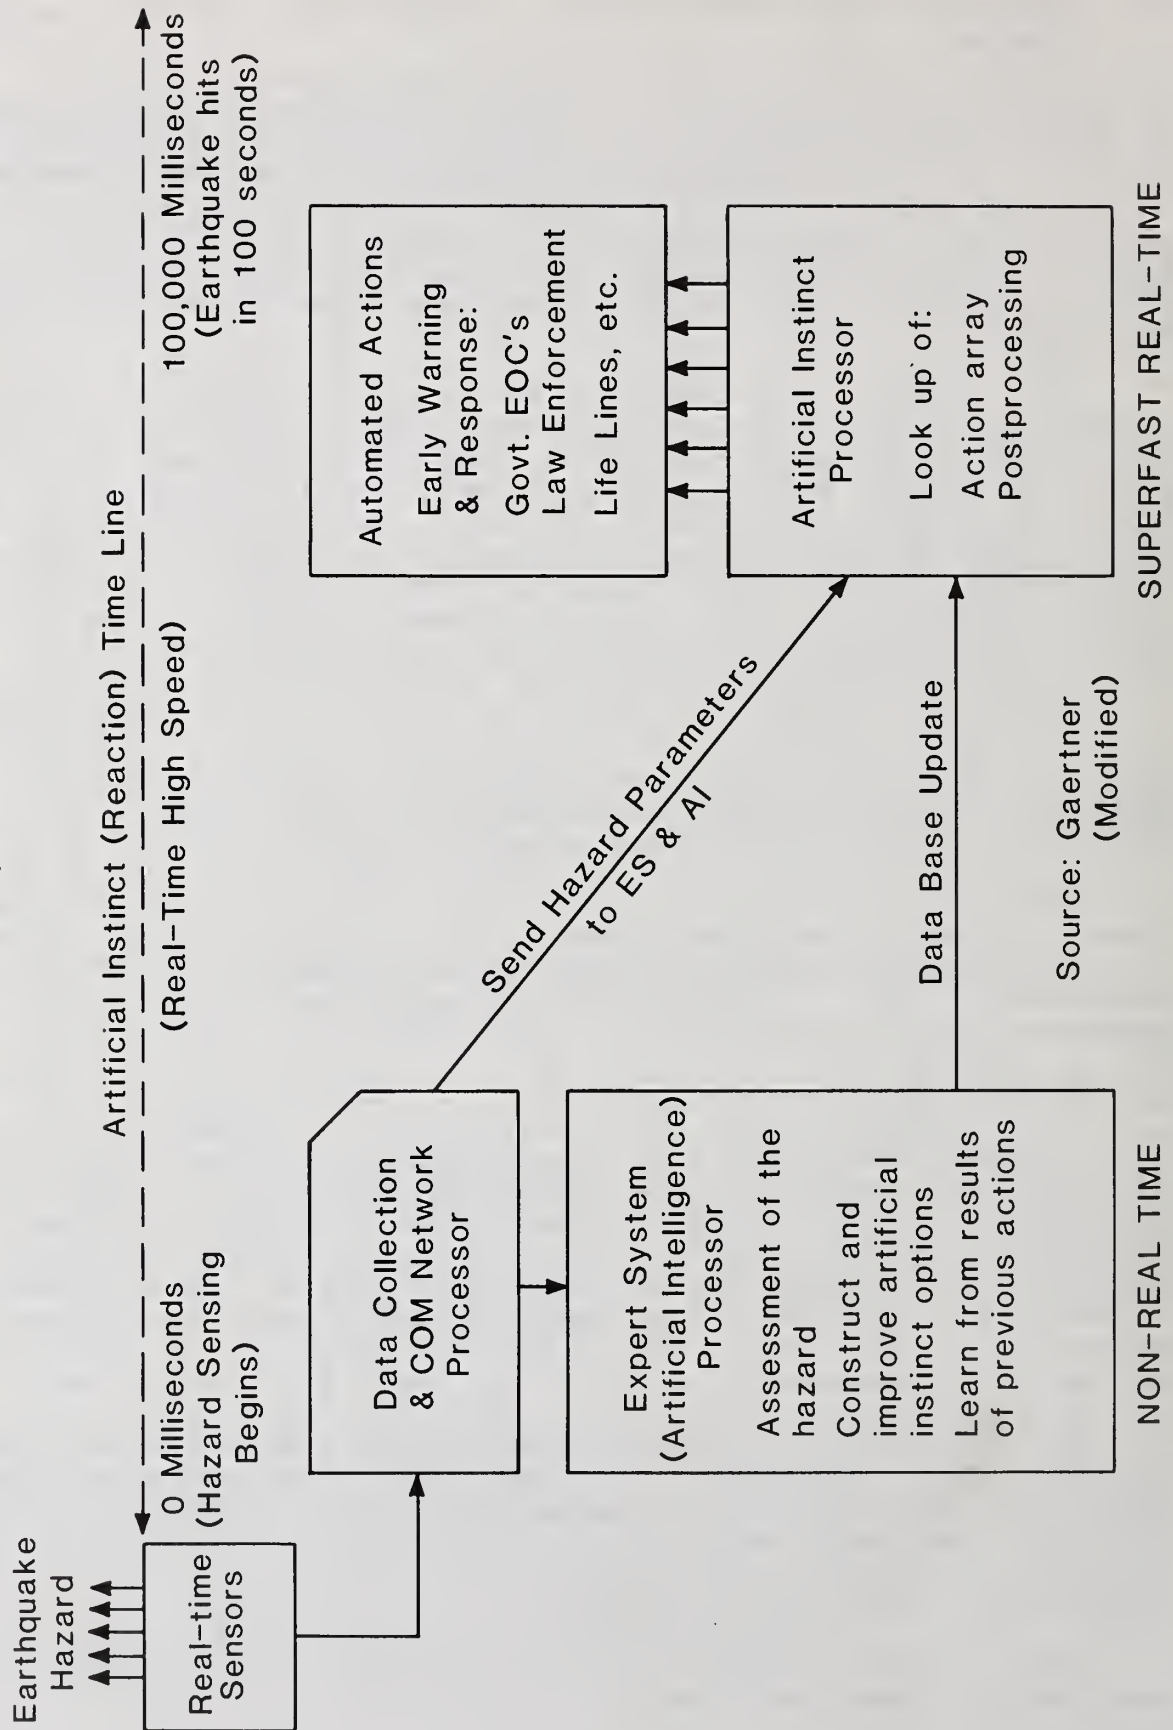
It consists of:

1. **KNOWLEDGE BASE** - the basic data, inference rules, and response procedures and frameworks associated with the particular environment. This artificial component of (AI)2 is considered "slow" or non-real-time because of the large size of the data base and the undesirability of attempting to process such a volume in a crisis where response must take place in a millisecond environment. This ES knowledge base, however, is the foundation for the next component;

2. **ARTIFICIAL INSTINCT MODULE** - this contains a super high speed processing capability for taking hazard condition indicators (search keys) and searching for an applicable set of prestored responses which are retrieved (there is a "hit") and then implemented for action in real-time to reduce the consequences of impending disaster.

The technical requirements for such a (AI)2 system are considerable (Gaertner, 1985):

Figure 6
ARTIFICIAL INSTINCT (AI)² ARCHITECTURE
FOR THE EARTHQUAKE PROBLEM



- specialized hardware to include fast circuitry with clock rates up to 100 MHz and large amounts of main memory (up to 64 Megabytes) for the artificial instinct component (Kimmel, 1985).

- a tested and reliable knowledge base with a good inference engine which can update both the knowledge base and, as a result, the artificial instinct data base as well;

- a reliable network of sensors to monitor the given environment;

- a fast, reliable telecommunications system linking various communities, organizations and levels of government via a wide area network (WAN) with linkages to other wide area nets (e.g., law enforcement, etc.) and key local area nets (LAN).

9. RECOMMENDED APPROACHES TO ES DEVELOPMENT IN EMERGENCY OPERATIONS.

The following are some suggestions and observations to the conference participants and FEMA about the potential advantages associated with various research options for ES prototyping and testing of concept.

9.1 DEVELOP ES PROTOTYPE AT FEMA TRAINING CENTER.

Use of the FEMA National Emergency Training Center as a base for development of a demonstration project in expert systems employing the training scenario of "Central City" which is currently the basis of exercises is highly desirable. Such a demonstration project (a prototype for training purposes) should be modular in scope and use simulated (non-real) data to develop quick, early ES versions which can be tested and evaluated by practitioner students who come for training to this center. Revisions can then be made and new versions issued for further evaluation. Widely available hardware and a suitable expert system generator software should be employed.

This approach would be similar to one being employed by the Defense Systems Management College (DSMC) which is developing its Program Management Support System (PMSS) which will have a knowledge base component. PMSS is designed to become the micro-based computer information system to be used by all major weapons acquisitions program management offices of the Department of Defense. The concept behind PMSS is modular, prototype development using private consulting firms. Both a training and a field prototype are to be developed. A similar approach should be taken by FEMA with regard to ES.

9.2 DEVELOP FIELD PROTOTYPES VIA DEMONSTRATION PROJECTS FOR SPECIFIC DISASTER TYPES.

Three possible applications to the emergency management area deserve FEMA research attention:

9.2.1 EARTHQUAKE WARNING AND RESPONSE.

Current discussions at the California Institute of Technology and other research institutions about the technical possibility of building a sensor-computer-communications network to give Southern California, for example, a 100 second warning of an imminent, disastrous earthquake clearly call for ES application. Significant response to such a short-lead time warning can only be done through plans requiring automated execution. ES would have a major role in determining not only whether a warning is justified but, after giving the warning, carrying out a series of actions (e.g., fail safe of commercial and other computers centers; securing of life lines such as transport, utility generation etc.) to reduce the destructive impact on the public and the infrastructure. This type of split-second response, if successfully carried out, could save numerous lives and reduce the necessary time for economic and other recovery by weeks if not months.

9.2.2 TOXIC CHEMICAL DISCHARGE WARNING AND RESPONSE.

This involves a series of complex models or a knowledge base designed to simulate/estimate the migration of specific hazardous materials on land, water, and other environments in a specific incident. This is likely to be the most difficult to develop since it is dependent on detailed data for a specific geographic area which in many cases is not available before an incident.

An ES component to such a system would be very valuable in those industrially vulnerable population centers where timely warning, containment and evacuation would not be otherwise possible.

9.2.3 NUCLEAR POWER PLANT ACCIDENT WARNING AND RESPONSE.

Work is already proceeding to integrate ES into the fail safe systems of this particular industry. In the response area especially, concurrent processing by ES based systems at several levels of government simultaneously would be especially desirable (Comfort, 1985).

Any program by FEMA to address the issue of how ES can be best applied to the emergency management area must include in its goals research in the (AI)2 area (de Balogh, 1985).

10. NATIONAL CLEARING HOUSE FOR TECHNOLOGY TRANSFER IN DECISION SUPPORT AND EXPERT SYSTEMS.

Finally, there is a definite need for a national repository of research findings and a clearing house for computer hardware and software information (including but not restricted to decision support and expert systems) in emergency management. It could provide software and even hardware at cost to prospective user agencies. The "computer store" concept now

used by the General Service Administration (GSA) is a good example.

An on-going program carried out through FEMA, via a number of universities with recognized research specializations in emergency management (e.g., by type of disaster) and significant library holdings on their particular topic is the answer.

11. SOME CONCLUDING OBSERVATIONS.

11.1 PROBLEMS WITH THE CONSORTIUM APPROACH TO ES DEVELOPMENT.

Since FEMA plans to proceed in ES evaluation primarily through work that the Department of Transportation will do with the recently acquired TIMM software from General Research Corporation, it is important to understand that this approach is not necessarily desirable or effective especially when development is by another organization.

To recall a bit of data processing history, the USAC experiment in MIS in the early 1970's had similar objectives and resulted in very little if any capacity building. Then the intent by a team of federal agencies with several million dollars to spend over several years was to develop a prototype MIS designed to be applied by any medium-sized American city. Both major functional (e.g., public finance, law enforcement, human resources, etc.) and integrated approaches to such development were undertaken with six cities being involved as test bed environments. This effort failed because it was technically not possible to transfer the results of that effort nor were the cities themselves interested in accepting concepts they had no hand in developing. This was a costly failure and an important lesson in technology transfer which must not be repeated with ES.

11.2 EMPLOYING USER FRIENDLY ES DEVELOPMENT TOOLS.

It should be noted that there are now other federal efforts underway to evaluate ES software which should be actively followed by FEMA. Foremost among them is the Department of Defense which has chosen the KEE (Knowledge Engineering Environment) software as its main ES developmental tool and has awarded Intellicorp, KEE's manufacturer, a million dollar contract to tailor KEE to DOD needs. KEE is a very user friendly system which merits close scrutiny by the emergency management community (Kinnucan, 1985; Univ. of Miami ICS, 1985). TIMM, on the other hand, does not have the same level of user friendliness.

A multiple research and evaluation approach conducted in parallel by FEMA using several excellent ES development (skeleton) software packages such as KEE, in several emergency contexts, is more appropriate. The time to find the right vehicle for capacity building, via ES, would then be considerably shortened.

11.3 PLAN NOW FOR TECHNOLOGY TRANSFER.

Based on our work at USC in decision support systems for the emergency management community, one conclusion is manifestly clear. Unless we systematically and purposively build into any expert systems research program a technology transfer mechanism which will build capacity of the emergency management infrastructure, we will not make a difference any time soon. Such a program must include: (1) clearly defined objectives and schedule for evaluation; (2) a workable organizational vehicle that will lead to information sharing within the emergency management community; (3) reliance on "ubiquitous" computer systems (hardware and software) that are widely available; and finally (4) a commitment to provide for technology transfer, even if that transfer meets only partial needs, NOW.

SELECTED REFERENCES

American Petroleum Institute "The Migration of Petroleum Products in Soil and Ground Water - Principles and Countermeasures," Publication No. 4149, (1972).

Bond, George "Clout and Salvo," BYTE, Volume 9, No.11, pp.279-286. Callamaras, Peter V. "Savvy," BYTE, Volume 9, No.2, pp.303-306. Duffy, Richard "CLOUT Narrows PC Language Gap," PC WEEK, Volume 1, No.17, pp.1,6. Eisenberg, Jane and Hill, Jeffrey "Using Natural-Language Systems on Personal Computers," BYTE, Volume 9, No.1, pp.226-238. Mace, Scott "Can Natural Language Sell?," InfoWorld, Volume 6, Issue 46, pp.36-41. Perrone, Giovanni "Clout," PCWeek, Volume 1, No.31, pp.74-75.

Campbell, Bruce. "FEMA: Information Processing in Emergency Preparedness", SIGNAL, Vol. 38, No. 11 (July 1984), pp. 35-37.

Comfort, Louise K. "Information Search Processes in Emergency Management: Computer Simulation as a Means of Improving Organizational Decision-Making Capacity," EMERGENCY PLANNING, San Diego, Ca.: Society for Computer Simulation, 1985, pp.41-45.

de Balogh, Frank. EXPERT SYSTEMS AND ARTIFICIAL INSTINCTS: A NEW CONCEPT FOR REAL-TIME WARNING AND RESPONSE FOR EMERGENCY MANAGEMENT, Los Angeles, Cal.: Decision Support Systems Laboratory, 1985 (In Preparation).

de Balogh, Frank. DECISION SUPPORT FOR OIL AND TOXIC CHEMICAL SPILL RESPONSE MANAGEMENT, Los Angeles, Ca.: Decision Support Systems Laboratory, January 1985 22p.

de Balogh, Frank. FUNCTIONAL AND TECHNICAL REQUIREMENTS FOR A PROTOTYPE EARTHQUAKE DECISION SUPPORT SYSTEM: A SURVEY REPORT, Los Angeles, Ca.: Decision Support Systems Laboratory, February 1985, 58p (Developed under NSF Grant CEE-8217406).

de Balogh, Frank. EXPERT SYSTEMS AND ARTIFICIAL INSTINCTS: A NEW CONCEPT FOR REAL-TIME WARNING AND RESPONSE FOR EMERGENCY MANAGEMENT, Los Angeles, California: Decision Support Systems Laboratory, University of California, DSSL-85-3 (In Preparation).

de Balogh, Frank, Petak, William, and Sessler, John. "Earthquake Decision Support System: A Computerized Demonstration Model for Earthquake Mitigation," *Proceedings of the First International Earthquake Conference*, Los Angeles, Ca.: February, 1985; pp. 273-293.

Foster, Edward. "Artificial Intelligence: Beyond the Buzzwords," *Personal Computing*, April 1985, pp. 62-69.

Gaertner, W.W. "(AI)2 For Real-Time." *Signal*, April 1985, pp. 55-61.

Giuffrida, Louis O. "FEMA Information Systems: Key to Effective Emergency Management," *SIGNAL*, Vo. 39, No. 9 (May 1985), pp. 147-156.

Hearne, Stephen, Hearne Laura, Mason, Sara, and Vesay, Anthony. TOXIC SPILL CONTAINMENT (USING LOTUS 1-2-3), Stuttgart, Germany: Institute of Safety and Systems Management, University of Southern California, 1984, Research Report.

Johnson, Carrol K. and Jordan, Sara. "Emergency Management of Inland Oil and Hazardous Chemical Spills: A Case Study in Knowledge Engineering". In, Hayes-Roth, Frederick (ed.), BUILDING EXPERT SYSTEMS. Reading, Mass.: Addison-Wesley, 1984. pp. 349-397.

Kimmel, Steven. "Using Very High Speed Integrated Circuits for Army Systems," *SIGNAL*, Vol. 39, No. 5 (January 1985), pp. 21-24.

Kinnucan, Paul. "Software Tools Speed Expert System Development," *High Technology*, March 1985, pp. 18-20.

Moore, Donald, Okamoto, Tatsuo, Russo, Joseph, Wilson, Robert; and Rojan, Christopher. "The FEMA Earthquake Damage and Loss Estimation System (FEDLOSS)," In, Carroll, John M. (ed.), EMERGENCY PLANNING, San Diego, Ca.: Society for Computer Simulation, 1985, pp. 118-124.

Petak, William (ed.). "Emergency Management: A Challenge for Public Administration." *PUBLIC ADMINISTRATION REVIEW*, Vol. 45, (Special Issue January 1985), pp. 172.

Rao, J. K. MODELS OF SEISMIC RISK AND DAMAGES FOR EARTHQUAKE DECISION SUPPORT SYSTEMS, Los Angeles, Ca.: ISSM, University of Southern California, June 1985 (forthcoming). (Part of the NSF funded EDSS project.)

Rojahn, C. and R.L. Sharpe, R.E. Scholl, A.S. Kiremidjian, R.V. Natt, R.R. Wilson. "Development of Earthquake Damage and Loss Estimates for Facilities in California," Eighth World Conference on Earthquake Engineering, Vol. 1 SEISMIC RISK AND HAZARD, Englewood, Cliffs, N.J.: Prentice-Hall, 1984, San Francisco. (Basis of unpublished 1984-FEMA/ATC-13 report).

Schuerman, L., Gardiner, P., Petak, W., and, de Balogh, F., "Decision Support System for Earthquake Mitigation and Response Management." (A proposal submitted to the National Science Foundation by U.S.C.'s Social Systems Research Institute [SSRI] and the Institute of Safety and Systems Management [ISSM], Los Angeles, California; July 1982). This contains a more detailed presentation of the EDSS concept.

Shah, H.C. and Dong, W.M. "A Re-Evaluation of the Current Seismic Hazard Assessment Methodologies," Eighth World Conference on Earthquake Engineering, Vol. 1 SEISMIC RISK AND HAZARD, Englewood Cliffs, N.J.: Prentice-Hall, 1984, San Francisco.

Sprague, Ralph and Carlson, Eric. BUILDING EFFECTIVE DECISION SUPPORT SYSTEMS, Englewood Cliffs, N.J.: Prentice-Hall, 1982.

University of Miami, Intelligent Computer System Research Institute. "DOD: \$1 Million to Intellicorp: Company Will Develop Expert System Tool Kit Based on KEE," ICS Applied Artificial Intelligence Reporter, Vo. 2, No. 7 (April 1985), p. 1, 5.

Wallace, William A. and de Balogh, Frank. "Introduction to Decision Support Systems in Disaster Management." Public Administration Review (January 1985) Special Issue on Disaster Management, pp. 134-146.

OPTIMUM EMERGENCY MANAGEMENT: THE EFFECTIVE USE OF INFORMATION TECHNOLOGY

Robert Lee Chartrand

Senior Specialist in Information
Policy and Technology

Congressional Research Service
Library of Congress*

The contribution that the emergency management community can expect from new information technology will not come from better or faster gadgets, but from greater understanding of the group dynamics of crisis situations, from a greater ability to prepare for such situations through sophisticated simulations and training, and from a finer analysis of their qualitative and quantitative characteristics.(1)

Dr. Jacques F. Vallee

Throughout countless generations mankind, from its origins in primeval forests, atop stilted lake dwellings, or in nomadic travail amid the inhospitable barrenness of desert regions, has coped with crises of its own making or caused by nature. The often murky chronicles of disparate civilizations tell the stories of such awesome happenings as the Great Inundation, the engulfment of Atlantis and Antilla, the slashing devastation wrought by massive meteors, and the inexorable glacial onslaughts. In many instances, the reportorial accuracy of the observer on the scene may be somewhat distorted, which once led Evelyn Waugh to note that we live in a world "in which perceptions count more than facts."(2)

Philosophically, our Nation prides itself on systematically analyzing situations and marshalling the requisite responsive resources--human, technological, budgetary--to plan for or cope with perceived or occurring crises. Information technology, through its spectrum of sophisticated mechanisms and methodologies, may offer an amplified context for those who must delineate a pragmatic philosophy of true emergency management. The intrinsic nature of human beings, acknowledged experts in some cases, reflects an awareness of the criticality of "stretching their horizons" which in this modern age may include the use of technology-supported outreach systems. This is indicative of their awareness that change is constant and requires a sustained effort on the part of institutions and individuals to adjust their infrastructures and goals.

There is a growing sense of urgency within both the public and private sectors that we as a Nation--at the Federal level as well as in the States and localities--remain inadequately

*The views expressed are those of the author and are not necessarily those of the Congressional Research Service nor the Library of Congress.

postured to cope with the spate of natural and technological crises of these complex times. Our societal structure and the governmental apparatuses conceived of over the years increasingly are hard-pressed to anticipate, much less respond effectively to, some of the disasters which have begun to impact lives and property. Surely those officials faced with the devastation wrought by the Mount St. Helens' eruption or on a lesser scale the scatter-shot tornadoes which ravaged the Carolinas early in 1984 would argue that a rethinking of certain priorities and processes should be a paramount consideration.

One critical facet of strengthening the emergency management capability in the United States--whether the focus is on mitigation, prevention, response, or recovery-- is the development of advanced information systems for collecting, storing, processing, retrieving, and sharing essential data and "value-added" information that may be used by emergency managers.

The introduction of advanced technology into any information handling environment is fraught with imponderables. The emergency management (EM) community is no exception. The growing panoply of sensors (aloft and aground), processors, and disseminators--the result of human inventiveness--instigates the rethinking of former, established positions. Included among the organizations in the United States with recognized roles and responsibilities in emergency management who share such concerns are:(2)

Federal agencies, such as FEMA, VA, USGS, NWS, NOAA, DOD, and many more; State and local governments (including task forces); regional commissions; private sector consultants and information services; information "clearinghouses"; organizational "watch centers"; national coordinating groups; and private sector contractors--corporate, university not-for profit.

Today, there is a sizable cadre--more than 12,000 designated "emergency managers"--responsible for anticipating and coping with crises of natural or man-caused origins ranging from tornadoes and floods to terrorism and toxic spills. These forces have learned that James Michener was correct when he pointed out that:(3)

...[our] balance in life consists of handling in real time those problems which cannot be delayed, then recalling more significant data during periods of reflection, when long-term decisions can be developed.

Congressional Insight and Initiatives

Congressional cognizance of this increasingly vital area has been on the upswing, and not surprisingly encompasses all four phases of emergency management, whether occurring before, during, or after a disaster. Among the many issues coming under initial legislative scrutiny are these of a higher priority.(4)

- * Is there a current, valid long-range plan addressing the role of communications networks in emergency situations?
- * Has the optimum use of advanced information technologies in various disaster scenarios been studied, and plans for their operational utilization developed?
- * Have priorities been determined for the creation, maintenance, and use of those essential information files which may be available to decision makers during emergencies?
- * Will secure communications be available in contingency situations?
- * Are the advantages and disadvantages of various technologies employed in anticipating or responding to natural or man-caused disasters understood by those managers or operators?
- * Is there a need to review present emergency management concepts and plans, particularly those concerning the roles of "watch centers," networks, and key human resources?

During the 97th and 98th Congresses, the Subcommittee on Investigations and Oversight, under the leadership of Representative Albert Gore, Jr., of the House Committee on Science and Technology, undertook a multi-faceted exploration of the role of information technology in emergency management. Early on, the Chairman noted that:(5)

The subject of disasters is not one that many of us care to dwell on. Earthquakes, fires, assassinations, terrorist attacks and nuclear melt-downs are the stuff of Hollywood and we like to keep it that way. As a result of this "out of sight, out of mind" ethic, our society is often ill-equipped to deal with emergencies when they do arise.

Present and potential uses of computers and telecommunications received primary attention in the 97th Congress opening series of events as the Subcommittee considered their value in preventing or coping with technological or natural

disasters. Following a roundtable discussion led by Representative Gore and Dr. Richard Beal of the White House, two days of hearings were held on September 29 and 30, 1981, with expert testimony by acknowledged governmental and private sector leaders in the field. Subsequently, a technical forum was sponsored by the Subcommittee on November 23, 1981, featuring participation by 17 senior individuals who engaged in structured discussions and witnessed demonstrations of technology-supported information systems.

The second phase of Subcommittee activity in this focal area occurred in the first session of the 98th Congress, with the convening of a combined two-day hearing and workshop, on November 16 and 17, 1983. The highlights of these various Subcommittee initiatives along with interpretive commentary and summary recommendations are contained in a committee print entitled "Information Technology for Emergency Management," prepared by the Congressional Research Service. The essence of the thematic goals enunciated and pursued by the Gore Subcommittee throughout its investigation is found in these four areas:(6)

1. a grasp of the full range of natural and technological disasters;
2. what technology can and cannot do;
3. the overt and subtle interaction between human beings and their innovations; and
4. the value of incremental improvements, when sweeping policy and program revamping is not possible.

A Range of Crises

Experience has shown that information is critical in anticipating or dealing with many types of disasters, and as improved technologies have become available they have been added to the inventory of equipment and related methodology which can be applied to EM tasks. A selective recital of disaster types and the havoc ensuing from their occurrence is instructive:(7)

- Boating and aircraft emergencies -- several hundred persons died in 1982 off-shore sinkings: and in FY 1982, 6,414 fishing vessels required SAR (search and rescue) assistance, plus 49,834 such incidents involving recreational boats. In addition, the FAA reported (in FY 1982) that 3,394 air carrier and general aviation accidents took place.
- Floods -- During 1983, according to the National Oceanic and Atmospheric Administration (NOAA) 204 persons died and property losses exceeded \$4 billion as a result of floods caused by accumulated snowfall runoff, dam breakage, or intensive rainfall.

- Hazardous materials spills (in transit) -- The National Transportation Safety Board (NTSB) estimates that 250,000 hazardous materials shipments are made daily. Transportation accidents--train, truck, barge--involving such cargoes have accounted for about 25 deaths per year over the past decade, in interstate accidents alone. In 1971 there were 2,225 hazardous releases, but by 1980 the statistic had increased to 16,115.
- Nuclear powerplant accidents -- Although Three Mile Island is the only "major" accident in recent years, there were 3,804 LEP (Licensee Event Reports) documented instances (in 1980) where performance exceeded technical design parameters. Filed with the Nuclear Regulatory Commission (NRC), these reports came from 69 commercial nuclear plants: 753 accidents were attributed to human errors and 2,174 to equipment malfunction.
- Tornadoes -- These destructive weather aberrations account for millions of dollars in damage and numerous deaths each year. The "super outbreak" in a 10 State area in 1974 involved 148 reported tornadoes which caused 315 deaths and \$600 million in property loss. In FY 1982, FEMA responded to 181 tornado warnings or touchdowns.

Other types of emergencies, while happening with less frequency, are still of high concern to Federal, State, and local officials alike. Earthquakes, volcanic eruptions, and hurricanes can be highly destructive and while infrequent when compared with some other forms of disasters, can wreak severe damage. Domestic terrorism, still modest in numbers, remains a menace of serious concern to authorities. In 1981, 31 terrorist incidents were reported in the United States, compared to 52 in 1982.

The Many Facets of Emergency Management

Even the most cursory examination of this responsible field of authority oversight and action provides an insight into the kaleidoscopic nature of emergency management which, it should be understood:

- occasionally brings about awesome but largely anonymous "miracles:
- has its share of impresarios, and sometimes such skills and talents are desperately needed
- often must be hurriedly established in "oases" of power, light, and other sustainers amid desolation
- usually is comprised of major episodes accompanied by singularly vital byplays

- may be represented, in more cases than usually imagined, by only the most embryonic "systems."

Decisionmakers at all levels of the emergency management structure have to understand the benefits and limitations of using information technology. Their questions include: What can technology do to better manage vital information? Can technology enhance the determination of information validity, including dealing with misinformation or "disinformation?" Should small-scale experiments and practical case studies be undertaken to better understand the man-machine interfaces involved in advanced systems?

A few important studies and related documentation have been produced in the past few years which represent a first look at this complex topic:

- 1979-1980 A series of papers on a proposed Crises and Emergency Management Information System (CEMIS), prepared by The MITRE Corporation(8)
- 1982 The Role of Science and Technology in Emergency Management, a project report by the National Research Council(9)
- 1982 FEMA Database Requirements Assessment and Resource Directory Model, a project report by the Information Retrieval Research Laboratory(10)
- 1983 Computer Simulation in Emergency Planning, conference proceedings(11)
- 1984 Remote Sensing and the Private Sector, a technical memorandum prepared by the Office of Technology Assessment(12)

Other means exist, in literature and the media, for broadening both professional and public understanding of disasters and existing or projected coping mechanisms. Veteran observers still recall, perhaps with bemusement, the reaction to Orson Welles' Martian invasion broadcast or the delight registered by those reading James Thurber's "The Day the Dam Broke."(13) Recent media offerings attracting widespread attention and comment have included such "docudramas" as "The Day After" (ABC, November 20, 1983), "The Crisis Game" (ABC, November 22-25, 1983), "Special Bulletin" (NBC, March 20, 1983), and "If I Were President" (ABC, August 6, 1980).

Retaining the Paramountcy of the Human Role

While most will aver that order and change need to be better synthesized within the emergency management environment, there are still compromises to be struck as to how such a meld is best achieved. Dr. Robert F. Kahn, long a moving force in his leadership role with the Defense Advanced Research Project Agency (DARPA), reminds us that "Technology for emergency management should be usable...a natural part of the working environment...dependable...ubiquitous...[and] interoperable with other systems."(14) This is a large order and may be seen by many at the local level, with limited resources, as unattainable. Yet, much has been done through the exercise of ingenuity, as in Vermillion County, Illinois, where a formidable emergency response force has been built on a "shoestring."(15)

There is a great deal that may be learned by civil sector authorities from their defense establishment and intelligence community counterparts. Time and again during the recent congressional hearings and workshop allusions were made to tools and techniques developed by these emergency-oriented groups. In many instances, it has been imperative to involve both military and civilian forces in responding to certain types of emergencies. It has been learned that not only is the sharing of resources during a disaster incident often required, but prior joint planning of such manpower and materials' disposition can be most beneficial. Another major contributor to coping with crises is the National Voluntary Organizations Active in Disaster (NVOAD) comprised of key private sector groups.

Collaborative effort is often impeded by bureaucratic obstacles or personality frictions, but the overriding need to mount interjurisdictional emergency response mechanisms--either informal or fully legalized--has resulted in impressively mounted action on a growing scale. Maximum communication between authorities at various levels must be achieved, not only to ensure protection of life and property, but because as Alexander M. Hunter has pointed out local emergency managers must "embark upon vulnerability analyses and a prioritization of hazard problems within their own communities....What happens if [they]...fail to take advantage of information that is available...to deal with the matter of saving lives and saving property?....they are looking at serious liability problems."(16)

Thus, with all of the incredible advances gained as a result of the inventiveness of man, the human factor remains paramount. Working to maintain the stability which underlies the general welfare, while striving concurrently to ensure sufficient emergency preparedness and plan perceptively for "continuity in government" under unsought in extremis conditions, the leadership's dilemma is very real as it strives to initiate a new

order. The reverberations of such change are manifest in the evolution of financial, commercial, and social service, as well as emergency preparedness, systems. Robert F. Littlejohn of the Pinkerton organization broadens our context of consideration regarding the human role by stressing that:(17)

Crisis management...is actually the antithesis of mismanagement and it is something that has to be really fine-tuned: identifying crisis issues; forecasting what an impact would be on a company politically, socially, and economically; and looking at what is the probability of the crisis taking place.

Issues and Opportunities

As noted earlier, human industry and ingenuity have combined to provide an ever-broadening array of technology-supported emergency handling capabilities. Illustrative of the major uses of these so-called "wizard machines" are:(18)

- 800 minicomputer warning systems in use throughout the country;
- a variety of airborne platforms such as the NASA U-2 and NOAA Flying Laboratories, along with sophisticated satellites featuring multi-sensor collection systems;
- the large masses of data stored in computerized or microfilm files (e.g., National Hurricane Center);
- the rapid retrieval of key data utilizing on-line access systems, by users located in emergency operations centers (EOCs), mobile units, or other remote (local EM) sites: and
- the varied communications conduits--landlines, airwaves, including satellite systems (e.g., Inmarsat)--for transmitting key data.

In the case of aerial data collection (infrared, radar, conventional photography), the possibilities for use in advance crisis condition detection are many. Arthur C. Lundahl and Dino A. Brugioni, pioneer practitioners in the arts of photographic interpretation and photogrammetry, offer the opinion that:(19)

The future portends even greater opportunities because imagery can now be digitized, and the combining of imagery interpretation expertise with computer technology...provides us with numerous innovative applications. The enormous volume of imagery-derived data now under computer control provides untold opportunities for utilization by analysts in Emergency Operation Centers.

One valuable resource for many emergency managers is geographical data which can be instrumental in presenting, and helping analyze, potential impact areas, transportation routes, and population placement. Dr. Jerome E. Dobson, drawing upon his experience at Oak Ridge in creating automated geographical systems, comments that the primary effect of this system would (20)

"not be to supplant human intelligence and decisionmaking, but rather to speed up many calculations and judgments that are already being made in emergency situations. The greatest advantage would be a better understanding of where the impacts can be expected to occur....this understanding would be shared through a common set of information among planners and decisionmakers at all levels of authority."

Another notable application of information technology, observes Dr. Marilyn C. Bracken, is that associated with hazardous waste management. In that connection, she emphasizes the "enormous potential" of its use: "Obtaining the best information possible--being able to assess the uncertainty and limitations surrounding the data, in a critical time frame--is the objective of all...involved in emergency management."(21)

These commentaries are indicative of the breadth and depth of this topic, and as the investigation by the Gore Subcommittee progressed, an identifiable plan of action emerged which could be undertaken by those responsible for emergency management in this country(22).

Firstly, the Federal Emergency Management Agency (FEMA) was urged to consider such actions as:

- Creation of a uniform disaster reporting system.
- Establishment of a National Assistance Program Index.
- Expansion of present orientation and training capabilities and programs, including multi-sensor work.
- Development of a permanent simulation capability, including models to exercise a wide variety of data.
- Preparation of a five-year plan, featuring interagency information handling capabilities.

Secondly, other Federal agencies, either in collaboration with FEMA or through unilateral action, should:

- Create a "core crisis management mechanism," preferably within the Executive Office of the President.

- Reexamine the present role of the Federal Communications Commission in emergency communications.
- Establish a focal point in a designated agency to study the application of technology to non-DOD problems.
- Direct Federal technology providers to review new information technologies for possible EM use.
- Undertake the establishment of a civil sector communications network with qualities of flexibility and durability.

Thirdly, State and local emergency handling authorities were encouraged to:

- Establish or upgrade mutual aid assistance agreements, with particular attention to technological interactive support.
- Standardize communications frequencies at least within an identifiable potential disaster area.
- Develop "liability" groundrules to clarify decisionmaking protocols and priorities.
- Optimize their EM capabilities in the interest of "self-sufficiency," a prerequisite for certain types of disasters.
- Begin to prepare that "secondary layer" of EM applicable statutes and regulations, to augment or serve in the absence of Federal laws.

Fourthly, the Congress through its committee structure, was asked to contemplate:

- Further revision of the Communications Act of 1934.
- Intensified oversight regarding a range of emergency management activity areas, both pre- and post-disaster.
- Creation of a "national emergency communications network."
- Establishment of a requirement for the "dual use of technology" in funding new procurements.
- Delineation of a new "clearinghouse" activity which would collect, store, and make available prioritized EM data.

This spectrum of recommendations goes far toward establishing a viable course of action for the foreseeable future. The sequence and speed with which these are accomplished

may be affected, in some instances, by factors beyond the control of the implementing institutions.

Effective Crisis Management: An Unflagging Need

Thus, it is difficult to foresee when this Nation will not have to confront Mencken's "great problems that heave and lather," but in the meantime every attempt must be made to attain a posture of preparedness and response that utilizes all possible resources. If a "systems approach" is a viable means to that end, then the lessons learned within the aerospace, military, and intelligence establishments must be heeded. The importance of this strategem was heralded in 1966 during the deliberations of the National Commission on Technology, Automation, and Economic Progress, which asserted that:(23)

In short, what a systems approach implies is comprehensive planning so that we can trace out the effects, progressive and regressive, of any set of choices and decisions upon all other relevant decisions.

Although there is an understandable and necessary emphasis on the ways in which officialdom collects, processes, and uses a great variety of information, the needs of the public-at-large for certain enlightenment--whether on television or radio, in written publications, or through lectures and seminars--also must be considered and met.

The durability and flexibility of emergency management systems are critical parameters for functional effectiveness. There is an upswing in simulations and actual exercising of certain crisis handling systems to test their execution of vital operations under stress. This constitutes one facet of "quality control" which can be a crucial underpinning for any such system. To many, if not most, emergency managers, the criterion is the delivery of needed information that is accurate, timely, comprehensive (where possible), and relevant to the challenge at hand. The "system" or "tool" may be incidental, and if its performance is unreliable--whether in terms of linking networks or simply retrieving a piece of key data--then the responsible office may opt to return to a simpler, more trustworthy capability.

The acquisition, verification, and transmittal of information has always been critical in meeting a variety of crises ranging from small-scale localized disasters to larger emergency situations affecting a wide geographic area. Information technology as never before has come to offer a wide range of potentials for enhancing the effectiveness of crisis organizations--both governmental institutions and designated authorities within the private domain--responsible for emergency warning and notification, situation assessment, decisionmaking during crises, and dissemination of essential information for responsive action.

If the challenges of the past are any indication of the pressures to be borne by the emergency management community in the future, those who shape key policies and direct the resultant programs must continue to seek out all potential resources to meet those forthcoming confrontations. And implicit in this challenge is the opportunity for optimum utilization of innovative technology which when combined with masterful group dynamics will ensure our survival and well-being as a Nation.

FOOTNOTES

- (1) U.S. Congress. House. Committee on Science and Technology. Subcommittee on Investigations and Oversight. Emergencymanagement information and technology. Hearings. 97th Congress. 1st Session. September 29, 30, 1981. Washington, D.C., U.S. Government Printing Office, p. 150. [Commentary by Dr. Jacques S. Vallee]
- (2) Chartrand, Robert L. Information Technology for Emergency Management: The Many Potentials. CRS Review. Congressional Research Service. U.S. Library of Congress (A.A. Hoehling, ed.) Vol. 5, #10. November/December 1984, p.23
- (3) Michener, James A. Space. New York, Random House, 1982. p. 622
- (4) U.S. Library of Congress. Congress Research Service. "Information Technology in Emergency Management." Pamphlet produced by Robert L. Chartrand, Clark F. Norton, Nancy R. Miller, and Madeline Seidner for hearings and work shops held by the Subcommittee on Investigations and Oversight. House Committee on Science and Technology, November 16-17, 1983. "Prologue" by Albert Gore, Jr.
- (5) U.S. Congress. House. Committee on Science and Technology. Subcommittee on Investigations and Oversight. Information Technology for Emergency Management. 98th Congress, 2nd Session. October 1984. Washington, D.C., U.S. Government Printing Office, p. 123
- (6) Ibid., p. 10
- (7) Chartrand, Robert L. Information for Emergency Management: The Many Potentials. CRS Review. Congressional Research Service. U.S. Library of Congress. (A.A. Hoehling, ed.) Vol. 5, #10, November/December 1984, p. 22
- (8) Stauch, B. Information needs of a Pilot crisis and emergency management information system (CEMIS). The MITRE Corporation [McLean, Virginia] July 1980. p. 59 and Janicik, E. Bounds for the pilot crisis and emergency management information system (CEMIS). The MITRE Corporation [McLean, Virginia] October 1979. p. 33
- (9) National Research Council Committee on Emergency Management. The role of science and technology in emergency management. National Academy Press. Washington, D.C., 1982, p. 90
- (10) Tenopir, Carol, and Martha E. Williams. FEMA data base requirements assessments and resource directory model. Washington, D.C., May 1982. p. 123

- (11) Society for Computer Simulation. Computer simulation for emergency planning. [Proceeding of conference, January 27-29, 1983 San Diego, California] Ed. by John M. Carroll. La Jolla, California. January 1983. p. 115
- (12) U.S. Office of Technology Assessment. Remote Sensing and the Private Sector: Issues for Discussion -- Technical Memorandum. Washington, D.C. March 1984. p. 41
- (13) Thurber, James. The Thurber Carnival. Franklin Center, Pennsylvania. The Franklin Library, 1980. p. 360-366
- (14) U.S. Congress. House. Committee on Science and Technology. Subcommittee on Investigations and Oversight. Emergency management information and technology. Hearings. 97th Congress 1st Session. September 29, 30, 1981. Washington, D.C., U.S. Government Printing Office, p. 180 [Testimony of Robert E. Kahn]
- (15) U.S. Congress House. Committee on Science and Technology. Subcommittee on Investigations and Oversight. Information technology for emergency management. 98th Congress. 2nd Session. October 1984. Washington, D.C., U.S. Government Printing Office, p. 345-356, 383-384
- (16) Ibid., p. 358. [Commentary by Alexander M. Hunter]
- (17) Littlejohn, Robert F. "Crisis Management: A Team Approach." Remarks during the AAAS Symposium session on "Information Technology for Emergency Management." New York, Sheraton Centre, May 27, 1984
- (18) U.S. Congress. House. Committee on Science and Technology. Subcommittee on Investigations and Oversight. Information technology for emergency management. 98th Congress. 2nd Session. October 1984. Washington, D.C., U.S. Government Printing Office, p. 5
- (19) Lundahl, Arthur C., and Dino A. Brugioni. Aerial Photography: World Class disaster fighter. An article to appear in the forthcoming issue [1985] on "Information Technology for Crisis Management" of The Information Society journal
- (20) U.S. Congress. House. Committee on Science and Technology. Subcommittee on Investigations and Oversight. Information technology for emergency management. 98th Congress. 2nd Session. October 1984. Washington, D.C., U.S. Government Printing Office, p. 320 [Testimony of Jerome E. Dobson]

- (21) Bracken, Marilyn F. "Chemical Emergency Response: Risk Assessment and Information Resources." Remarks during the AAAS Symposium session on "Information Technology for Emergency Management." New York. Sheraton Centre. May 27, 1984
- (22) U.S. Congress. House. Committee on Science and Technology. Subcommittee on Investigations and Oversight. Information technology for emergency management. 98th Congress. 2nd Session. October 1984. Washington, D.C., U.S. Government Printing Office, 1984. p XV-XVI
- (23) Technology and the American Economy. Report of the National Commission on Technology, Automation, and Economic Progress. Vol 1. February 1966. p. 100

ROLES OF SIMULATION IN THE APPLICATION OF EXPERT SYSTEMS IN EMERGENCY MANAGEMENT OPERATIONS

by A. Ben Clymer
Clymer Technology
Ocean, New Jersey 07712

1. INTRODUCTORY CONCEPTS

1.1 Concepts Used

In this paper an approach is used that is rather different from the usual. The following assumptions have been applied:

1. The bionics approach to artificial intelligent systems design can be productive (i.e., drawing upon how various concepts of the intelligence of organisms, including man, are organized).
2. Block diagrams are a fruitful way to communicate many of the concepts.
3. The human intellect is highly parallel, not just serial like a classical Von Neumann computer.
4. It is necessary to deal with levels of abstraction and hierarchies of meaning, in order to understand some of the fundamental intelligent functions on both the afferent (input) and the efferent (output) side.
5. Analogies enable a further exploitation of the bionics approach, if a number of systems analogous to the system of concern are known.

1.2 An Approach to Human and Artificial Intelligence

The perceptual side of the central nervous system may be conceived to be essentially hierarchies of recognition devices, with at least one hierarchy of each sensory modality. Each recognition device ("recognizer") has an output at the top only when its set of inputs at the bottom meets certain criteria that correspond to some concept. The output is more abstract than the inputs. A hierarchy of recognizers can recognize very abstract concepts. See Figure 1-1.

The inputs to a recognizer need not be "on" (i.e., high or signaling) all of the time in order for a recognition to be achieved. Any uncertainty shown by inputs that are on intermittently (like blinking lights) will be shown in the same manner and to the appropriate degree also in the output. Thus, fuzziness in the inputs produces fuzziness also in the output. The higher the percentage of the time an output is on, the greater the certainty of its recognition.




FIG. 1-1: RECOGNIZER




FIG. 1-2: ELABORATOR

Some of the visual and tactile recognizer hierarchies are organized to identify patterns in a perceived field of two dimensions. It is conceivable that this process takes place by pattern matching, in which the system finds the figures that have to be synthesized and superposed in order to match the perceptual field. If there is a residual difference between the synthesized and perceived fields, figures are sought in the residual field and then are matched, until the residual field is a uniform or random background. This process requires that each synthesized figure be changed in size, location and orientation until it cancels out the corresponding figure in the perceived field.

A set of recognizers need not be a strict hierarchy (like an organization chart in which each person has only one boss). A recognized concept can flow upward to two or more recognizers in parallel. In principle there need be no fewer outputs at the top than there are inputs at the bottom. A recognizer can have more than one abstract concept in its top layer. In fact, as in some optical illusions, the concept recognized can alternate between two possibilities.

One of the benefits of having a field broken down into a small number of figures is that they can be stored much more efficiently than if every bit in the field had to be stored. For recall of the field it would be necessary only to call out the figures and superpose them.

On the neuromuscular side of the central nervous system there are hierarchies of command elaboration units ("elaborators"), each of which converts a single input into a number of distinct outputs. The input is more abstract or general than the outputs. See Figure 1-2.

A hierarchy of recognizers is itself a recognizer. Therefore, the graphic symbol for a recognizer (Figure 1-1) can be used to denote either a recognizer or a hierarchy of recognizers. This convention can save a lot of detail in diagrams. See Figure 1-3.


Similarly, a hierarchy of elaborators is itself an elaborator. The graphic symbol in Figure 1-2 can be used for either. See Figure 1-4.

Another component of the central nervous system is the arousal system (reticular system). It dictates the functional level that must be aroused and involved at the present moment; higher levels can be essentially at rest. It takes its input from the hierarchies of recognizers for all sensory modalities. It sends an arousal signal upward, and it sends commands for increased gain coefficients in all recognizers and elaborators that are appropriate to the arousal.

Yet another component of the central nervous system is a reflex unit. It has one or more inputs from recognizers, and it has an output to an elaborator. Thus, it can cause an action appropriate to the recognition, which is a reflex action.

There are also lower level components, such as control (homeostatic) loops, which adjust physiological variables when there are changes in


RECOGNIZERS ARE BUILDING BLOCKS.
THEY CAN BE PUT TOGETHER IN A HIERARCHY:



THIS HAS A TRIANGULAR ENVELOPE, SO A
RECOGNIZER SYMBOL CAN BE USED ALSO
TO REPRESENT A HIERARCHY OF
RECOGNIZERS.

FIG. 1-3: RECOGNIZER HIERARCHY

ELABORATORS ALSO ARE BUILDING BLOCKS.
THEY TOO CAN BE PUT TOGETHER IN A HIERARCHY:



THIS TOO HAS A TRIANGULAR ENVELOPE, SO AN
ELABORATOR SYMBOL CAN BE USED TO
REPRESENT A HIERARCHY OF ELABORATORS.

FIG. 1-4: ELABORATOR HIERARCHY

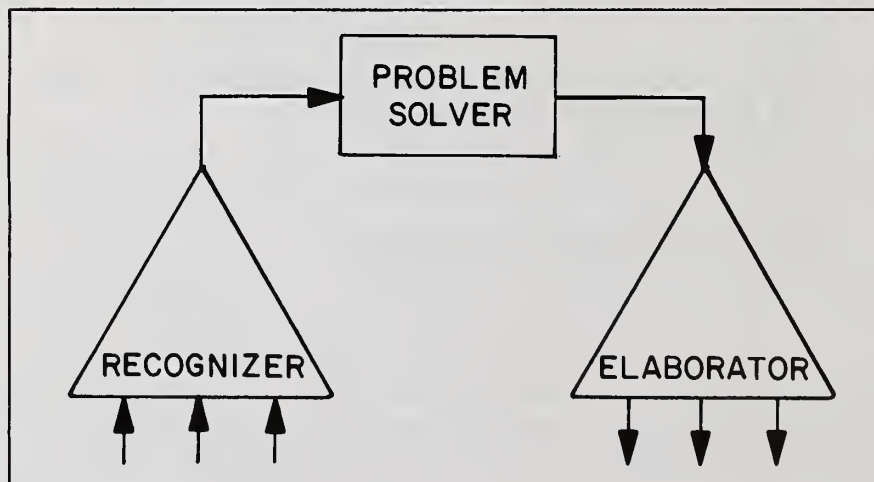


FIG. 1-5: INTELLIGENT SYSTEMS

the body's boundary conditions. These lower levels are not of concern in artificial intelligence.

All of the foregoing components working together cannot duplicate the capabilities of the human brain (or any other animals' brains). One of the missing capabilities is problem solving, planning, decision making, etc., which can be dealt with only partially in today's artificial intelligent systems. These activities are among a manager's functions.

One of the key functions of a manager is planning. A way to conceptualize planning is in terms of pattern recognition in an abstract sense, using some appropriate and familiar conceptual scheme to provide structure. For example, the plan needed might be an internal portion of a Pert chart, given the rest of the chart. As in pattern recognition, building blocks would have to be "seen" that together would fill the gap and meet the requirements, then directing attention to any requirements not met. Some heuristic algorithms for tasks of this type are well known in artificial intelligence.

1.3 Some Systems Analogous to the Nervous System

These systems have analogous structures:

1. A business manager.
2. Expert system.
3. Battlefield officer.
4. Management support system.
5. A full-blown intelligent robot.
6. A comprehensive AI system.
7. A hierarchy of managers and/or expert systems in any mix.
8. A hierarchy of agencies in emergency management.

These systems have an analogy also with the central nervous system. It might help in thinking about these systems to portray them all as shown in Figure 1-5.

However, there are some distinct differences between the nervous system and an Emergency Management (EM) system. For example, in the EM system:

1. There are no clearcut hierarchies of recognizers and elaborators. The elements of the system, mainly people, have both in each person.
2. There is no specialized arousal system in many EM systems.
3. The EM systems do not have characteristic signal frequencies like the beta, alpha, gamma, and delta rhythms.
4. There is very little proprioception.
5. There is some sensorimotor action (as in redirecting the direction of gaze), such as when a person is sent to a particular area to inspect and report.

6. Uncertainty in EM messages is not manifested in the transmission.

1.4 Computer Implementations

Some of the functional systems discussed above can be rendered with various kinds of computer hardware. The implementations with LISP programs and AI serial digital computers are well known and so will not be discussed here.

Parallel computers have aroused some strong support recently in artificial intelligence circles (see References 1 and 2), partially as a result of the Japanese national computer project. Generally, there is a trend in artificial intelligence to look more closely at any commercial parallel computer or parallel arrays of computers or both in one system. One rationale is to increase speed; another is to emulate how the brain is believed to work, which is surely largely in parallel. One purpose is to simulate an expert system, and another purpose is to have an expert system.

It might be helpful conceptually to think of the parallel type of computer processing repetitively a hierarchy of recognitions, row by row, from bottom to top, starting over as often as input changes warrant. One could think of each recognizer that is actually recognizing something as lighting up. The pattern of lights would show what is being recognized. Similarly, a hierarchy of elaborators would show what elaborations it is producing in this pattern of lights. Blinking lights would signify uncertainty.

Another type of implementation that is of interest in artificial intelligence is a robot. Conventional robots are programmed to perform repetitively a useful motion in an open-loop manner (no proprioceptive feedback). A hierarchy of elaborators could, in principle, generate the detailed motions appropriate to a general command, yet be able to respond to other commands also.

2. SIMULATION IN EMERGENCY MANAGEMENT

2.1 Meanings of "Simulation"

"Simulation" has come to have several meanings in the EM field. One meaning is that of a game or exercise, played for training, research, or plan development, having participants who play roles of emergency management personnel according to a preestablished scenario. This meaning is sometimes subdivided according to whether a computer is used in an important role. With a computer involved, one has a "technology-supported simulation," as opposed to a "traditional simulation" if a computer is not involved (Reference 3).

A totally different meaning of "simulation" is taken from the field of simulation, namely, the capability or activity of generating, by means of a computer, the behavior of a system by progressively solving in time

the dynamic equations comprising a mathematical model of the system. Often this is called "computer simulation," a term which, unfortunately, can mean the simulation of a computer, which is an important branch of simulation. As applied in EM, this meaning of "simulation" is the simulation of an emergency itself, usually without real or simulated individual people. However, as discussed in Section 2.5, simulation of people in emergencies is under development and will appear more and more in EM simulations.

A "simulation," in the sense of a capability of a computer program, generates an approximate version of the simuland's (of the system to be simulated) behavior in specified situations and for specified purposes. Cases outside the region of the specified situations and purposes are usually not rendered with much realism.

2.2 Functions of Simulation for EM

Simulation has valuable uses in all four phases of EM (mitigation, preparation, response, and recovery):

o Mitigation

1. Simulations of emergencies and mitigation efforts to determine the best allocation of resources among emergencies of different kinds to get the best return on the EM investment by reducing the incidence and/or impact of certain emergencies.
2. Detailed simulations of a particular class of emergency to determine for it the best allocation of resources among mitigation activities that might be effective, and to evaluate the cost/benefit ratio.
3. Simulations of compound emergency causation, hazard probability, vulnerability, etc., for local situations such as chemical or power plants.

o Preparation

1. By means of simulations of emergencies and responses to them, it is possible to determine the best plans for given contingencies. Some of these studies could well be done by local industry, which has a real concern for emergencies.
2. Simulation is a powerful tool for R&D on emergency responses by the various people involved, particularly in matters of psychology and group dynamics. Real people would be used in this research.
3. Simulation of EM management processes with real people included is potentially valuable for testing particular organizations and for evaluating types of organizations, particularly under stress of a realistic exercise.

4. Simulation is the tool of choice for most studies in system engineering to improve the performance of systems for communication, information processing, etc.
5. Simulation of certain emergencies for prediction of where and when the threat is to be expected, for the purpose of short-term planning and blending into the response phase.
6. Simulations can be used to generate dramatic demonstrations of emergencies and the consequences of various public responses, good and bad, for public education presentations.
7. Simulations of all types can be used effectively in training professionals in EM or affording them practice. The whole continuum of degree of stress is available for use in particular cases.

o Response

1. Simulation can be used as a manager's analytical tool in support of decisions during an emergency, such as to project the future severity of the emergency, including the effects of the planned response.
2. Simulation can be used to determine unknown parameters of an emergency (such as release quantity or rate, or wind speed and wind direction averaged over a plume) by matching iteratively the predicted plume to the field data concerning the plume.

o Recovery

1. Simulation can be used to compare the benefits of alternative plans for recovery.
2. Simulation is often used as a tool for reconstructing what must have happened in an emergency, step by step, matching recorded information, such as to ascribe the cause, and learn whatever is possible for next time.

The foregoing applications of simulation to EM would normally be performed locally for local purposes. However, the national EM organization has echelons, each of which has reason to do its own simulation studies. For example, on the national level one could single out:

1. Simulation of the human ecology of the country, with emphasis on its vulnerabilities and risks (Reference 3).
2. A national facility for EM simulations of all kinds on a regional or national scale (Reference 3).

2.3 Simulation of Emergencies: State of the Art

There has been fairly general agreement in EM on the scope of the meaning of "emergency." However, in recent years there has been arising a good case for extending the time scale of "emergency" to include very slow processes that pose severe threats to man.

Examples would be the threat to the ozone layer by aerosols, the greenhouse effect due to increased production of carbon dioxide, the opposite effect due to processes acting in the other direction, the progressive exhaustion of forests and other resources, acid rain, the theorized poisoning of the Romans by lead, the multiple consequences of exceeding "the limits to growth," etc., etc. Many of these slow emergencies have been simulated already. Some papers of this type will be presented at the Conference on Simulation at the Frontiers of Science, sponsored by the Society for Computer Simulation, Norfolk, Virginia, March 1986.

Even without including these slow emergencies, there is a considerable body of past simulations of emergencies. The longest bibliography known to the writer is in Reference 4. This bibliography lists 45 fairly recent papers under the following headings:

1. Simulation of Fires and Firefighters.
2. Simulation of Emergencies in Plants.
3. Simulation of Spills and Decontamination.
4. Simulation of Structural Emergencies.
5. Simulation for Water Management.
6. Simulation of Deployment of Specialized Vehicles.
7. Simulation of Emergency Medical Services.
8. Simulation of Evacuation of Buildings and Cities.

The frequency distribution of publication dates indicates a steady growth in the number of papers per year.

Many applications and potential applications were absent from the foregoing bibliography: wind storms over land or water, ocean waves and surges, avalanches and mudslides, nuclear weapon effects, storage tank accidents, outdoor fires, vehicle fires, building bombings, civil defense systems, and chained multiple emergencies (e.g., an earthquake breaking a dam, which floods its plain, or a storm affecting transportation and communications). Most of the omissions were due to the selectivity in the data base that was searched.

The Proceedings of the Conferences on Computer Simulation in Emergency Planning, sponsored by the Society for Computer Simulation, San Diego, January 1983, and January 1985, include some of the areas omitted above.

A bibliography on evacuation research has been compiled by John Sorenson, of Oak Ridge National Laboratory. Some of these papers report on the use of simulation as a tool for evacuation research.

2.4 Simulators for EM

2.4.1 Introduction

Definitions and meanings of the word "simulator" were the subject of an entire panel discussion at the SCS Simulators Conference, Norfolk, Virginia, March 1985. Roughly speaking, however, for the present purpose, a "simulator" will be a special-purpose simulation capability implemented on a dedicated computer. It might or might not have a videodisk or graphics as an adjunct.

Here some actual and some possible future simulators are presented to illustrate what simulators can do for EM. Since simulators form a \$2 billion per year industry which is eyeing EM as a market, emergency management leaders might well now look seriously at simulators as potentially valuable tools.

2.4.2 Puff and Plume Prediction

The largest and most advanced facility for simulating the transport and spread of a plume or puff of radioactive pollutant is ARAC (Atmospheric Release Advisory Capability), at Lawrence Livermore National Laboratory (References 5 and 6). It provides calculated consequences of releases of radioactivity into the atmosphere, such as level of surface contamination or cumulative dose at a point, resulting from a plant or transportation accident. Current and predicted values are available in tens of minutes from Lawrence Livermore National Laboratory. Over 50 facilities of DOE and DOD are served by ARAC.

The simulation programs used by ARAC cover a wide range of needs. In a site system at each facility of concern there is a personal computer which uses a Gaussian plume model for prompt information to EM personnel. The site computer is also tied to the main computer at LLNL, where the model used (ADPIC) is 3D particle-in-cell fluid transport and diffusion. The wind field, which is consistent with the actual detailed topography of the site, is calculated by the MATHEW program and is used as input to ADPIC. The central computer also provides graphic displays of isolines for the plume, dose, and surface contamination. It is probably only a matter of time until ARAC will be extended to include toxic chemical or biological releases, volcanic dust clouds, etc.

A recent survey (Reference 7) of computer systems for assessing plumes turned up quite a few that are on the market: the Safer System, Chemical Emergency Evaluation System and Community Assurance Program, Hazardous Assessment System for Toxic Emissions, Chemical Hazardous Area Release Model and EMERGE.

The least well known parameter is the magnitude of the release expressed as a release mass or a release rate, depending on the accident. It is possible to determine the release magnitude or rate with the puff or plume, respectively, simulation by iteratively adjusting it to give the best fit to whatever downwind measurements of pollutant are available.

The same principle can be extended to determine source location or wind parameters, if they are unknown or not known accurately. Figure 2-1 shows the principle.

2.4.3 Flood Prediction Systems

The same principle can be applied to the flooding of a river. Rain gages and water depth gages in the tributaries report automatically to a central computer, which is able to interpolate any missing data and then solve the flow equations to forecast flood levels for the entire river basin. The technology has already been demonstrated by International Hydrological Services, and it was reported at the Emergency 85 conference in Washington, D.C., in May 1985. It was predicted that there will be 1,000 automated flood computer systems in the U.S. by the year 2000.

2.4.4 Nuclear Power Plant Simulators and Analyzers

For the past 20 years there have been, all over the world, an increasing number of multi-million-dollar simulators of nuclear power plants for training operators. The trend is for all but the smallest of plants to have their own simulator. These simulators can give training in a wide range of situations arising from malfunctions or operator errors, including severe emergencies. However, in general, the simulators are unable to provide training in a situation requiring action by more than one or two operators in the control room. Only their own actions are noted by the simulator. This is one future direction of evolution of simulators.

The Three Mile Island incident stimulated the development of another class of simulators called "analyzers," which are capable of running a plant simulation faster than real time for research, for on-the-spot planning, or for later reconstruction of an accident. Analyzers are only now becoming able to meet the stringent requirements on speed of calculation of a huge model.

2.4.5 Outdoor Fire Simulators

In fighting a large outdoor fire, it is valuable to have a portable simulator for planning. The U.S. government funded development of a simulator called "BEHAVE" for this purpose (Reference 8), among others.

2.4.6 Some Future Simulators

One can imagine a dozen or more distinct types of simulators that would be valuable in EM. Their time will come.

For example, one elaborate type of simulator, which has been described by Clymer and Perley (Reference 9), would contain three levels of the overall EM organization associated with a nuclear power plant: the control room operators, the utility management, and the outside agencies. Emphasis of the uses (R&D, planning, and training) would be upon the processes of interaction among the three organizational echelons.

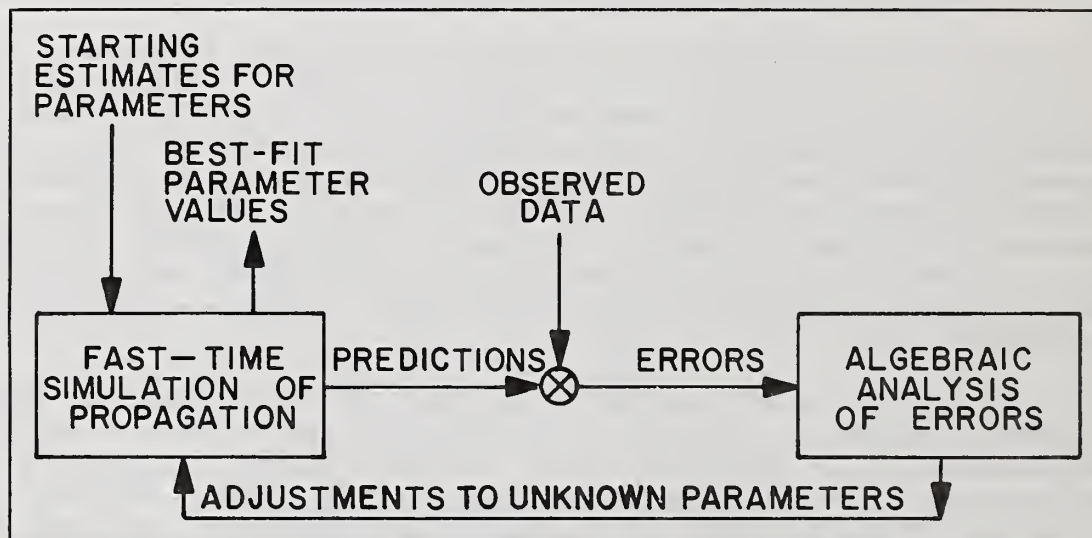


FIG. 2-1: ITERATIVE DETERMINATION OF RELEASE PARAMETER VALUES

Therefore, the simulator need not include all of the details of each person's environment or task.

Another type of simulator would be highly generic but flexible. That is, it would be designed to represent any of a set of plants. The control panels would be replaced by generic CRT displays and keyboard commands. The software would enable the user to construct a new plant by drawing the flow diagram on the CRT screen, the software producing the appropriate equations for the model from a modular data base. Such a simulator would be useful for a state or regional agency concerned with several plants.

The other end of the market is a large number of agencies having small budgets. Thus, there is an opportunity for development of simulators on personal computers. Among the possible uses would be planning and training. Each agency would simulate mainly its own operations, plus simplified interactions with superior, inferior, and lateral organizations in the EM hierarchy.

2.5 Modeling People for EM Simulations

At present there are three distinct approaches to the modeling of people, such as those in EM organizations and citizens in an emergency. The oldest was continuous dynamic modeling of variables in simple situations, such as a single homeostatic loop in physiology or the tracking of a target by a gunner. The second has been the use of logic and discrete simulation technology in a digital computer for prediction of sequences of events, such as the motion of a pilot's line of sight over an instrument panel in flight situations. The third is the use of expert systems to stand in place of people for restricted expert functions. These three approaches are illustrated below with some examples.

The continuous dynamic approach gives a transfer function in the case of gunner or pilot, or a set of differential equations in the case of a physiological system. These models do not reach very far upward in the human functional hierarchy, but they do provide a foundation. Continuous models succeed significantly in describing these systems, which themselves have a continuous character.

There are other systems, however, which are apparently not continuous but which offer promise of being modeled by differential equations nevertheless. An example is the "position" of a hostage while captured by terrorists. The intermittent and sudden changes of position are similar to the motions of a block of wood on a tabletop with static friction and two opposed springs pulling on it. One spring represents the force applied by the hostage's traditional beliefs, and the other spring represents the force applied by the terrorists. This mechanical model can be described "exactly" by differential equations. On the whole, however, systems involving beliefs, feelings, decisions, etc., are not able to be modeled well by differential equations.

Discrete models are preferable for accounting for a series of discrete purposeful actions. There are several types of human actions that might be or have been modeled this way:

1. What an operator of a plant will do in a particular situation could be modeled in terms of three, say, actions that might be taken, each with a given probability each time. Two of the three would be "human error," each with a small probability. The probabilities could be made functions of the stress level.
2. A similar model might be applied to members of the public in an emergency in following evacuation directions and road signs, or trying to dial a phone.
3. Likewise, the model might apply to many aspects of the work of an emergency manager.
4. Another type of model deals with a human trying to keep track of many variables, such as a pilot scanning an instrument panel or plant operator scanning a whole control room. Depending on the previous rate of change of a variable, its proximity to a critical value, and the variance of the variable, he will return to a certain variable within a time period estimated to be safe, attending to more urgent variables meanwhile. Work of this kind was pioneered at Bolt Beranek and Newman.
5. When a human operator is to be modeled in a task network, the SAINT language has been used (Reference 10).

The third approach, using artificial intelligence to simulate some human ability, has been evolving rapidly. The first major breakthrough was symbol processing for problem solving (Newell, Shaw and Simon). Recently, the most promise has been shown by expert systems in situations requiring complicated decisions, diagnoses, etc., which will be discussed herein in Section 3. An alternative approach has been to apply fuzzy logic in simulating human controllers, such as operators forced to go into the manual mode of control systems in an emergency. See References 11 and 12. It is interesting to note that the fuzzy controller tries to satisfy each of several conflicting guiding principles, compromising the deviations, a process that is reminiscent of and partially analogous to the foregoing spring-and-mass model of a hostage's position.

All in all, human simulation has a long way to go, but some of the roads are passable. Sorenson, of Oak Ridge National Laboratory, has stated that the best theories and models of behavior account for only 40 to 50% of the variance (Reference 13).

3. EXPERT SYSTEMS IN EMERGENCY MANAGEMENT

The desirability of applying expert systems to emergency management was recognized by Frederick Hayes-Roth in Reference 3. The first major

public exploration of the possibilities has been the First Symposium on the Application of Expert Systems in Emergency Management Operations, of which this paper and these Proceedings are fruits.

In this section a brief treatment is given of the topic of the Symposium, not so much for its own sake as for providing entree to the subject promised by the title of the paper, which is treated in Section 4.

3.1 A Concept of Expert Systems

Expert Systems as a discipline is considered to be a subset of Knowledge Engineering, which is a subset of Artificial Intelligence (AI), Reference 14. The functions of Knowledge Engineering Systems include problem solving, search, and symbolic programming, and these subsume the functions of expert systems (see Section 3.3).

Jay Liebowitz has defined expert systems as "computer programs that mimic the behavior of a human expert in a specific domain of knowledge." Thus, expert systems are simulations in a sense. However, they bring a new class of models to the simulation field.

Expert systems have some sort of logic in their core, where inferences are produced. The outermost layer is a natural language interface for bilateral communication between the spoken language of the user and the formal language in the core.

The core, where knowledge is stored, data are brought in, and inferences are made, can be strictly logical and deterministic. However, it is also possible for an expert system to work with fuzzy logic and to generate probabilistic or equivalent semi-quantitative expressions of natural language.

Another parameter in the meaning of "expert system" is the degree of expertise. Some expert systems exceed the ability of a human expert, some match it, but even less ability would be noteworthy and useful, particularly early in the development of a new application area. It would not be appropriate to apply the term "expert system" to a much less impressive degree of expertise, such as the models of people that are typical of the field of human simulation today (see Section 2.5), or an automated reference librarian or file clerk.

3.2 State of the Art of Expert Systems

The telling and retelling of the successes achieved by expert systems has been the chief fuel rocketing artificial intelligence into prominence. Here only a brief summary is given.

The spearhead function for expert systems, with which the first major breakthroughs were made starting in the 1960's, is diagnosis and troubleshooting of disease, malfunctioning hardware or software, etc. Many other functions and applications have followed, as is discussed in Sections 3.3 and 3.4. However, Dixon (Reference 15) argues that only

about four expert systems could be considered operational in 1983, although the number has been growing. The first two expert systems to equal and then surpass human expert capability were Dendral and Macsyma (Reference 14). Another well publicized success has been the system to plan the manufacturing of VAX computers of varied configurations.

Expert systems are growing also in complexity. The current upper limit on program size is 1,000 to 2,000 rules.

Emergency management is only lately becoming a recipient of the benefits of expert systems. Among the current developments is an expert system designed by Radian Corporation to assist in making decisions during emergencies (Reference 7).

3.3 Functions of Expert Systems

3.3.1 Input Data Processing

The functions of expert systems are classified here in terms of three categories:

1. Input data processing
2. General data processing
3. Output data processing

The first category consists of convergent-upward hierarchical data processing, as described for recognizers in Section 1. Similarly, the third category consists of divergent-downward hierarchical data processing, as described for elaborators in Section 1. The second category consists of processing that includes both the first and third but also characteristically a very general set of operations such as problem solving, designing, planning, etc.

Input data processing involves six or more media in which the input data might be expressed, as shown in Table 3.3-1. In each medium the processing is most often and perhaps most successfully arranged hierarchically, as suggested by the column headings of the table. The analogies among these operations in different media are recognized in the generic row at the bottom of the table which encompasses them all.

The process shown in the first row of Table 3.3-1 is called pattern recognition, machine vision, computer vision, automated perception, etc. An expert system for object recognition has been the subject of a paper (Reference 16). The same process includes inspection of manufactured parts for defects. A recent review covers today's technology (Reference 17).

The second row of the table addresses a similar but more complicated case in which the perceiver is moving relative to the objects in the field of view. Before the visual field can be analyzed it must first be rid of the effects of the observer's motion, by determining what that motion is. Then the process proceeds as in the static image case.

Table 3.3-1
Input Data Hierarchical Processing

	<u>Information Medium</u>	<u>Preliminary Processing</u>	<u>Progressive Recognition</u>	<u>Top-Level Recognition</u>
1.	Static Image	Edge Detection, Image Processing, Enhancement	2½D Pattern Recog- nition	3D Object Rec- ognition by Matching
2.	Moving View	Moving Edge De- tection, Self- Motion Identifi- cation and Correc- tion	Pattern Recog- nition	Object Recog- nition
3.	Numerical Data	Data Processing and Storage/ Retrieval	Parameter Severity Identification	Situation Recognition
4.	Natural Written Language	Word-Level Recognition	Message/Passage Interpretation	Gist Recog- nition
5.	Medical Language	Symptoms Identification	Syndrome Deter- mination	Diagnosis of Disease
6.	Natural Spoken Language	Frequency and/or Analyses of Phonemes or Words	Word-Level Pattern Rec- ognition	Gist Recog- nition
7.	Any or All of the Above Data	Intermedium Translation and Merge	Abstract Pattern Rec- ognition	Concept Recognition

The medium in the third row is numerical data. An example could be a situation the severity of which is being assessed by frequently measuring and communicating some variables of the situation. A system to monitor these data and interpret them progressively in a hierarchical manner for severity would be used also to store the data in a data base and retrieve any needed data. This would be an expert system in the management and analysis of data about a situation. For example, the situation could be a fire that has been "put out" in a building, and the expert system could be left in the building at a key location to watch for signs of the fire breaking out again.

Written language, the medium of the fourth row, subsumes all processing of natural language in written form as input data. The desired outcome can be identification of the gist of a message, or it could be an assessment of a situation described in words.

A special case is diagnosis of a disease described in medical language, as shown in the fifth row of Table 3.3-1. The processing deals progressively with more general sets of symptoms, culminating in the syndrome of the disease finally diagnosed. It is not true that all diagnosis systems work this way, but they could. An analogous problem is diagnosis of a fault in an electrical or mechanical system, given a set of facts in written language.

The corresponding problem stated in spoken language is shown in the sixth row. It too involves analyses and recognition of progressively larger units of thought. A recent review of these systems has been given in Reference 18. The present art is compelled to make compromises such as requiring a limited vocabulary, or limiting how many different speakers can be dealt with, or forcing the user to separate his words with pauses, or otherwise falling short of the desired capabilities. Expert systems are expected to have a major role in upgrading speech recognition system performance.

Since all of the foregoing media applied similar hierarchies of progressive processing, it is possible, as shown in the seventh row, to generalize to the processing of any given mixture of the media in a hierarchy of successively larger recognitions.

Nothing has been said specifically about emergency management applications of these recognition systems. Yet applications are possible with hierarchical processing in any of the media shown in Table 3.3-1.

3.3.2 Problem-Solving Expert Systems

The term "problem solving" is intended here to embrace all human intellectual effort except pure recognition and pure elaboration. This category includes designing and planning as well as recreational or practical solving of problems of any type. In Figure 3-1, a problem solver block is shown acting in a closed loop with the outside world, making use of a recognizer and an elaborator, to solve a problem.

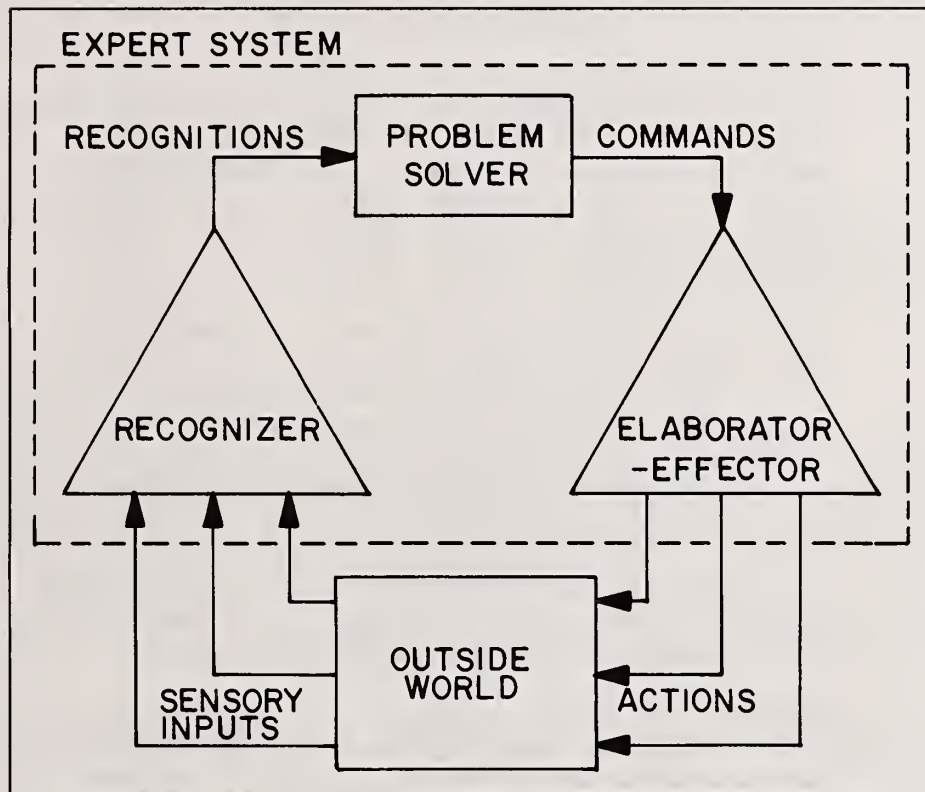


FIG. 3-1: CONFIGURATION OF EXPERT SYSTEM
INTERACTING WITH OUTSIDE WORLD

A rather artificial classification of problem solving systems is made here into "knowledge systems" and "action systems." In the special case of knowledge systems, the only outside world of concern is a computer or specialized tool for knowledge processing, and its output is knowledge. The output of an action system is action directed at the outside world in the unrestricted sense.

The following examples of expert system functions that can be classified as knowledge systems are given in approximate increasing order of difficulty:

1. An artificial instructor, such as a computer program and interactive video, could be developed to teach a student a specified body of knowledge. In the simplest case the only feedback from the student is a map of knowledge and ignorance as he progresses. (Question answering is deferred to Item 6 of this list.)
2. An expert system during an emergency can manage on-line analyses, predictive simulations, and probabilistic forecasts of the future developments in an emergency. These services on behalf of a decision maker can significantly increase the effectiveness of decisions relating to emergencies. Simulations can be run by an expert system also between emergencies for such purposes as research, training, and long-range planning.
3. A somewhat similar capability is to reason about a causality network, such as to determine the probable origin of a fire from indirect information from remote sensors.
4. Another function that could be undertaken by a knowledge system is machine learning or knowledge acquisition. Today's expert systems are given their rules in advance.
5. Expert systems are being developed to do mechanical design (Reference 15 and references therein).
6. Another function of a knowledge system is development of a plan, which is an iterative process that ultimately leads to a plan meeting all requirements. Simulation might function as a step in the design loop.
7. An expert system can give computerized support to an emergency manager in connection with decision making. These services include consulting, advice giving, question answering, etc. It is a short step from there to actual decision making in lieu of a person, such as in early or remote aspects of emergencies when emergency managers are scarce.

Action systems have much more interaction with the outside world than do the foregoing examples of knowledge systems. Several types can be distinguished:

1. Troubleshooting and diagnosis, on engineered systems or human bodies, alternating with tests whose outcomes divide the alternatives roughly in half (as in "20 Questions"), are such functions.
2. An area in which there has been intensive development is automatic picking of parts from a randomly loaded bin (Reference 19). Both machine vision and part gripping are involved.
3. An even more complicated function is autonomous remote fighting of an emergency under conditions too demanding for a human. Examples of such emergencies would be a radioactive spill to be picked up and removed or a fire generating very high temperature.
4. Still more complicated would be backup replacement of an emergency manager by an expert system, including the functions of decision making (within some limited domain) and crisis management.

3.3.3 Output Data Processing

The output data processing functions are pure command elaboration in a hierarchy driving a battery of effectors (muscles in the nervous system paradigm). As with input data processing there are several distinct modalities (media) in which the commands can be rendered: spoken language, computer code, graphics, written words, etc.

The following list of examples is indicative of the great potential variety:

1. Elaboration of alphanumeric graphic commands to a printer.
2. Progressive elaboration of computer code into more detailed languages, even for expert systems.
3. Progressive allocation of resources or deployment of entities to agencies of smaller scope.
4. Elaboration of repair instructions.
5. Automatic speech generation from symbolic words.
6. Progressive elaboration of commands down to a number of parallel orders or messages, to coordinate emergency fighters.
7. Hierarchical organization of a program elaboration function for a generic robot (instead of humans writing code directly for the lowest level of detail).

3.4 Applications of Expert Systems

Each application of expert systems is a function of expert systems as implemented in a particular field. Thus, expert systems constitute a large matrix. Such a matrix is given in Reference 20. Only a brief illustrative list is included here:

- o Equipment Operation and Maintenance
 - 1. Decision support for well drilling when problems are encountered (Reference 14).
 - 2. Equipment failure diagnosis (Reference 21), e.g., for locomotives, turbines, programmable controllers, computers, jet engines, cars, and telephone cables.
 - 3. System debugging (Reference 22).
- o Medical Diagnosis
 - 1. INTERNIST and CADUCEUS (Reference 23).
 - 2. Mycin, Emycin, Puff, etc. (Reference 14).
- o Industrial Plant Control
 - 1. Several companies are developing expert advisors for nuclear power plant operators (References 23 and 24).
 - 2. Expert system for control system design (Reference 25).
 - 3. Expert system for monitoring and interpretation of a multiplicity of alarms (References 25 and 26).
 - 4. Expert system for process control (Reference 27).
- o Vehicle Control
 - 1. Navigation by an autonomous vehicle (Reference 28).
 - 2. Expert system for emergency procedures for an aircraft (Reference 29).
 - 3. Pilot's assistant.
- o Automated Programming
 - 1. Simulation programming tools (Reference 24).
 - 2. Expert system for program writing (Reference 30).

- o Military
 1. Battlefield command decisionmaking support.
 2. Surveillance of ships with automated analytic aids (Reference 31).
- o Financial
 1. Legal tax advice.
 2. Fuzzy rules and logic for stock market advice.
- o Emergency Management
 1. Structural damage assessment (Reference 32).
 2. Risk management analyses.
 3. Fire detection and assessment (Reference 33).

The foregoing list is, as mentioned above, only indicative of the variety of expert system applications today. Some additional commercial expert systems are listed in Reference 34.

3.5 Development and Implementation Methodology

Expert system methodology has been evolving from its beginning in the LISP language and the Von Neumann serial digital computer. The evolution has led to quite a few distinct methodologies.

One of the branch points has been away from Boolean logic into fuzzy logic. "Reveal" can be used to program an expert system in fuzzy logic (Reference 35). Reveal is a decision support program.

Another branch point is the method used to obtain the rules that underlie an expert's thinking in his domain of expertise. The traditional method is to interview an expert interminably until the rules are complete. Recently there has appeared an alternative method which uses a program written for a personal computer (Reference 36).

Expert system architectures are in ferment. The classical architecture consisted of a data base for rules and data, an inference engine, and perhaps a natural language interface. There are now several strategies implemented in inference engines (Reference 14).

There is a ferment also in the approaches being used for expert system development. The conventional computational logic approach, which processes "what ifs" in LISP, on a serial computer, is finding growing competition from approaches which are naturally suited to the speedier parallel computers (Reference 37). Kurzweil feels that "...parallel architectures are the wave of the future" (Reference 37).

The concepts in Section 1, such as recognizers, lend themselves to the parallel approach. A recognizer may be regarded as a forward inference

engine operating bottom-up on a hierarchy of rules to recognize a hierarchy of concepts. If implemented on a parallel computer, a recognizer would be traversed a row at a time at the computer cycle rate.

Some of the other approaches that are currently being developed or used are:

1. Decision theory, combining some ideas from statistics, operations research and economics.
2. Using simulation as a stepping stone (Reference 38).
3. Factor correlation as a method for determining domain rules from data for particular cases, avoiding rule programming.

One of the explosive trends is the appearance of a diversity of software packages that facilitate or obviate the programming of rules for expert systems. Many of these new software tools have been written for personal computers. The following list is a sampling:

1. GCLISP (Golden Common LISP) is an IBM PC implementation of LISP. (See critique in Reference 39).
2. Knowledge Engineering Environment (KEE), from IntelliCorp, is one commercial "shell" which saves programming time (Reference 40).
3. Personal Consultant, from Texas Instruments, based on the EMYCIN shell, runs on the TI personal computer (Reference 40).
4. M.1, from Teknowledge, runs on the IBM PC (Reference 40).
5. Insight Knowledge System, from Level 5 Research, also runs on the IBM PC (Reference 40).
6. SMALL-X (Reference 41).
7. Exsys, from Exsys, is for creation of expert systems on an IBM PC (Reference 42).
8. Expert Choice, from Decision Support Software, Inc., runs on the IBM PC or PC XT (Reference 43).
9. Expert Ease, from Human Edge Software Corp., also runs on the IBM PC or XT (References 36 and 44). It lets the user build his/her own rules by learning from experience.
10. Expert Systems Development Tool (ESDT), being developed for DARPA by IntelliCorp, is based on KEE.

11. The Intelligent Machine Model (TIMM), from General Research Corp.'s Advanced Technologies Division, runs on an IBM PC. It asks questions and reconciles the answers in order to obtain rules (Reference 45).

Most of what has been written about expert systems implicitly assumes that they will last forever. However, there will always be need for maintenance, upgrading, and updating, so the life cycle viewpoint is appropriate.

3.6 Future Evolution of Expert Systems

Expert systems will evolve in many dimensions such as:

1. New applications, e.g., management of calculations, and EM applications.
2. Improvements in performance by orders of magnitude, as for example in the areas of speech recognition and language translation.
3. Enormous growth in scale of effort, as illustrated by the 1984 to 1989 projected market growth in artificial intelligence from \$147 million to \$1.1 billion (Reference 46).
4. Expansion into larger and faster computers, which will accommodate more elaborate programs in expert systems.
5. Combinations of expert systems into more complex and capable systems, not only in integrated products but also in the separate but interactive mode of organization, leading to hierarchies of expert systems.
6. More flexible relationships between expert systems and people.

4. SIMULATION AND EXPERT SYSTEMS IN EM

The union of the fields of simulation and artificial intelligence has been called "Intelligent Simulation" (Reference 47). It is the main subject of Section 4, although some particular known applications to EM are mentioned.

4.1 Simulation in an Expert System

Simulation can be useful when it is a block within an expert system, as shown with two examples in Figure 4-1. To be sure, simulations are not normally parts of expert systems, but such a relationship is a possibility. One such case has been mentioned in Reference 48, which involves a simulation within an expert system.

Another possibility is an iterative organization of an expert system for design, in which one block could be a simulation of the performance of




FIG. 4-1: SIMULATION IN AN EXPERT SYSTEM

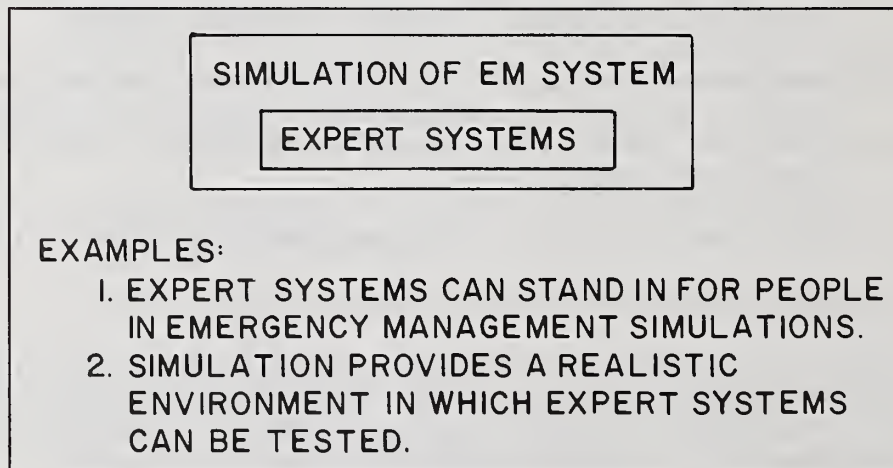


FIG. 4-2: EXPERT SYSTEM IN A SIMULATION

the system being designed. Also in future one could expect to see human operator or emergency manager models inside an expert system. In fact an entire EM hierarchy might be simulated within a multiagency EM expert system.

4.2 Expert System in a Simulation

Figure 4-2 shows examples of the opposite case, in which an expert system is one block inside a simulation. The first example is based on the fact that any expert system is, in a way and to a degree, a discrete simulation of a thinking person or team of persons or agency in EM, for instance. In Figure 4-3 there is shown an expert system used as a high level part of a simulation of a person. In the second example of Figure 4-2 a simulation of a domain provides an interactive environment for testing and development of an expert system.

An expert system could appear also inside simulation software for decisionmaking or planning for a model or run campaign. The simulation could be related to EM.

Expert systems are entering simulation. As put by Oren, "The question is not whether or not to have artificial intelligence in simulation, but rather how to have it? at which level? how reliably? how soon? and above all how intelligently?" (Reference 49).

4.3 Other Relationships of Simulation and Expert Systems

Rather than within or encompassing a simulation, an expert system can be over a simulation in a managerial sense. In an emergency, or in case there are to be thousands of runs (another type of emergency), it is valuable to have some automation.

Simulation has been found to be a good tool for use in evaluating and researching EM decision support systems (Reference 55). Similarly, as shown in Figure 4-4, simulation has been demonstrated to be a useful stepping stone in the development of an expert system for plant control (Reference 36).

Expert systems have been proposed or applied for the design of a computer simulation (References 51 and 50). Also, artificial intelligence has been examined for applicability to development of a complete simulation environment (Reference 52).

Simulation of computer cognition has been discussed (Reference 53).

The ROSS language has been developed for defining simulations, such as of large combat situations (Reference 54).

4.4 Future Developments

Simulation in relation to expert systems or artificial intelligence is the subject of several conferences in 1985, so it is a very active area

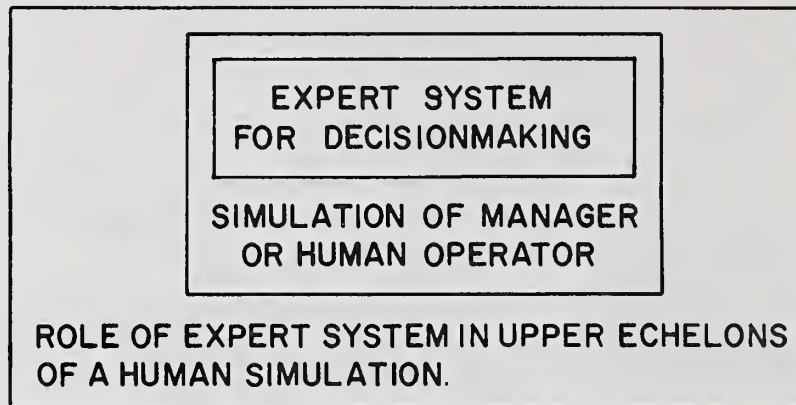


FIG. 4-3: EXPERT SYSTEM IN A
HUMAN SIMULATION




FIG. 4-4: SIMULATION AS A STEPPING STONE
TO AN EXPERT SYSTEM

of research. One of those conferences that is not mentioned in the References section of this paper is the conference on Artificial Intelligence in Simulation, University of Ghent, Belgium, February 25-27, 1985.

Reference 51 is indicative of the far-reaching thinking that has already been done toward further harnessing AI in the service of simulation. One sees less about the benefits that simulation will offer to artificial intelligence, but it should not be underestimated.

REFERENCES

1. Waldrop, M. Mitchell, "Artificial Intelligence in Parallel," *Science*, Volume 225, 10 August 1984, pp 608-610.
2. Feldman, Jerome A., "Connections," *Byte*, April 1985, pp 277-284.
3. Congressional Research Service, Library of Congress, "Information Technology for Emergency Management," U.S. Government Printing Office, Washington, D.C., October 9, 1984.
4. Clymer, A. Ben, "Simulation in Emergency Management in the Year 2000," Proc. International Conference and Exhibition for Emergency, Disaster Preparedness and Relief (Emergency 85), Washington, D.C., May 21-24, 1985.
5. Dickerson, M.H., et al., "ARAC Status Report: 1985," Lawrence Livermore National Laboratory, University of California, Livermore, California, Report UCRL-53641, May 1985.
6. Sullivan, Thomas J., "Modeling, Simulation and Emergency Response," Proc. 1985 SCS Conference on Computer Simulation in Emergency Planning, January 1985, San Diego.
7. Basta, Nicholas, "U.S. CPI Start to Feel Effects of Bhopal Tragedy," *Chemical Engineering*, March 18, 1985, pp 27-33.
8. Andrews, Patricia L., "A System for Predicting the Behavior of Forest and Range Fires," Proc. 1983 SCS Conference on Computer Simulation in Emergency Planning, San Diego, January 1983, pp 75-78.
9. Clymer, A. Ben, and Perley, Daniel, "Simulators for the Coordination of Nuclear Power Plant On-Site and Off-Site Radioactivity Emergencies," Proc. SCS Conference on Computer Simulation in Emergency Planning, San Diego, California, January 1983.
10. Laughery, K. Ronald, "Modeling Human Operators on a Microcomputer: A Micro Version of SAINT," *Simulation*, January 1985, pp 10-16.
11. Zadeh, Lotfi A., "Making Computers Think Like People," *IEEE Spectrum*, August 1984, pp 26-32.
12. Zadeh, Lotfi A., "A Fuzzy-Set-Theoretic Approach to the Compositionality of Meaning: Propositions, Dispositions and Canonical Forms," *Journal of Semantics*, December 1983, Volume II, No. 3/4, pp 253-272.
13. Sorensen, John, "Public Perspectives on Nuclear Power," *Emergency 85*, Washington, D.C., 21-24 May 1985.
14. Hayes-Roth, Frederick, "The Knowledge Based Expert System: A Tutorial," *Computer*, September 1984, pp 11-28.

REFERENCES (Continued)

15. Dixon, J.R., and Simmons, M.K., "Computers that Design: Expert Systems for Mechanical Engineers," Computers in Mechanical Engineering, November 1983, pp 10-18.
16. Kim, J.H., Payton, D.W., Olin, K.E., and Tseng, D.Y., Hughes Aircraft Company, "An Expert System for Object Recognition in Natural Scenes," First Conference on Artificial Intelligence Applications, Denver, Colorado, December 5-7, 1984, sponsored by IEEE Computer Society and American Association for AI.
17. Aldersey-Williams, Hugh, "Machine Vision Finds New Applications," Digital Design, June 1985, pp 40-46.
18. Meng, Brita, "Speech Recognition: Not a Typical Engineering Program," Digital Design, June 1985, pp 49-57.
19. Edson, Daniel, "Bin-Picking Robots Punch In," High Technology, June 1984, pp 57-60.
20. Frenkel, Karen A., "Toward Automating the Software-Development Cycle," Communications of the ACM, Volume 28, Number 6, June 1985, pp 578-589.
21. Herrod, Richard A., and Papas, Barbara, "Artificial Intelligence Moves into Industrial and Process Control," I&CS, March 1985, pp 45-49.
22. Gupta, N.K., and Seviaora, R.E., "Expert System Approach to Real Time System Debugging," First Conference on Artificial Intelligence Applications, Denver, Colorado, December 5-7, 1984.
23. Anon., "Artificial Intelligence," Science, Volume 223, February 24, 1984, pp 802-804.
24. Murphy, Thomas E., "Setting Up an Expert System," I&CS, March 1985, pp 54-60.
25. Foxboro, "Expert System Solves Control Problems," I&CS, March 1985, p 53.
26. Andow, Peter, "Alarm Systems and Alarm Analysis," Plant/Operations Progress, Volume 4, Number 2, April 1985, pp 116-119.
27. Moore, R.L, et al., "A Real-Time Expert System for Process Control," First Conference on Artificial Intelligence Applications, Denver, Colorado, December 5-7, 1984.
28. Parodi, A.M., "A Route Planning System for an Autonomous Vehicle," First Conference on Artificial Intelligence Applications, Denver, Colorado, December 5-7, 1984.

REFERENCES (Continued)

29. Lineberry, M., et al., "Intelligent Automation of Emergency Procedures in Advanced Fighter Aircraft," First Conference on Artificial Intelligence Applications, Denver, Colorado, December 5-7, 1984.
30. Haradvala, Sam, et al., "Expert Systems for High Quality Code Generation," First Conference on Artificial Intelligence Applications, Denver, Colorado, December 5-7, 1984.
31. Groundwater, Elizabeth H., "A Demonstration of an Ocean Surveillance Information Fusion Expert System," Proc. Artificial Intelligence and Simulation Conference, Norfolk, Virginia, March 1985.
32. Ogawa, H., et al., "An Expert System for Damage Assessment of Existing Structure," First Conference on Artificial Intelligence Applications, Denver, Colorado, December 5-7, 1984.
33. Carbonell, Jaime, "Sensor-Based Diagnosis in the SMOKEY System," these Proceedings.
34. Anon., "I&CS Guide to Artificial Intelligence Products," I&CS, March 1985, pp 50-52.
35. Okuma, Amy, "Software and 'Fuzzy' Logic Let Any Good Programmer Design an Expert System," Electronic Design, April 4, 1985, pp 173-184.
36. Crabb, Don, "Expert Ease 1.11 Reconsidered," InfoWorld, April 1, 1985, pp 46-50.
37. Kurzweil, Raymond, "What Is Artificial Intelligence Anyway?," American Scientist, Volume 73, May-June 1985, pp 258-264.
38. Elmaghraby, A.S., Jagannathan, V., and Ralston, P., "An Expert System for Chemical Process Control," Proc. AI and Simulation Conference, Society for Computer Simulation, Norfolk, Virginia, March 1985.
39. Amsterdam, Jonathan, "GCLISP," Popular Computing, May 1985, pp 116-118.
40. Kinnucan, Paul, "Software Tools Speed Expert System Development," High Technology, March 1985, pp 16-20.
41. Kaplan, Randy M., "SMALL-X: An Environment for Constructing Expert Systems on a Microcomputer," ACM Sigsmall Symposium on Small Systems, May 1985.
42. Heite, Ned, "Exsys," InfoWorld, January 28, 1985, pp 43-44.
43. Beechhold, Henry F., "Expert Choice," InfoWorld, January 28, 1985, pp 45-50.

REFERENCES (Continued)

44. Milman, Jeffrey O., "Do-It-Yourself Expert Systems with Artificial Intelligence on a Microcomputer," Proc. IEEE Compcon '84 Fall Conference on the Small Computer (R)Evolution, Arlington, Virginia, September 16-20, 1984, pp 20-26.
45. Elphick, Art, "TIMM--Expert System Builder," Simulation, June 1985, p 306.
46. Anon., "Artificial Intelligence: Current and Future Commercialization," Business Communications Co., Inc., Stamford, Connecticut, December 1984.
47. Anon., "Call for Papers," 1984 Conference on Intelligent Systems and Machines, Oakland University, Rochester, Michigan, April 24-25, 1984.
48. Kohler, J.L., "The VENTILATION EXPERT for Underground Coal Mines," these Proceedings.
49. Oren, Tuncer I., "Intelligence in Simulation," Simuletter, ACM Sigsim, January 1985, Volume 16, Number 1, p 3.
50. Vansteenkiste, Ghislain, summary of paper given at conference on Frontiers in Simulation Hardware and Software, Scottsdale, Arizona, April 11-12, 1985.
51. Shannon, Robert E., Mayer, Richard, and Adelsberger, Heimo H., "Expert Systems and Simulation," Simulation, June 1985, pp 275-284.
52. Reilly, Kevin D., Jones, Warren T., and Dey, Pradip, "The Simulation Environment Concept Artificial Intelligence Perspectives," Proc. AI and Simulation Conf., Norfolk, Virginia, Holmes, Willard M., Ed., March 1985.
53. Clancy, William J., and Shortliffe, Edward H., Eds., "Readings in Medical Artificial Intelligence: The First Decade," Addison Wesley, 1984.
54. Klahr, Phil, report on paper presented to Fourth UKSC Conference on Computer Simulation, Bath, England, Simulation, February 1985, p 97.
55. Bellardo, Salvatore, Karwan, Kirk R., and Wallace, William, "An Investigation of System Design Considerations for Emergency Management Decision Support," IEEE Trans. on Systems, Man, and Cybernetics, Volume SMC-14, Number 6, November/December 1984, pp 795-804.

Expert System for Fire Emergency Management

Geoffrey N. Berlin
Modeling Systems, Incorporated
1718 Peachtree Street
Atlanta, Georgia 30309

Emergency management is a highly inter-disciplinary field concerned with the arrangements and functional capabilities available to mitigate the injury and property damage due to natural and man-made disasters. The purpose of this paper is to discuss the application of an "expert systems" formulation to fire emergency management and to illustrate how local, state and federal officials can utilize such a capability to react to a variety of fire emergencies.

The proper management of emergency resources to combat fire can prevent thousands of injuries and millions of dollars in property damage. A typical fire grows at an ever-increasing rate in intensity and area. Since time is critical and the available resources are limited, the rapid deployment of firefighting resources is essential.

Fire certainly poses a major challenge to local, state, regional and federal emergency agencies. Exhibit 1 illustrates the number of major conflagrations that have occurred within the past eighty years. Fires have been major problems following military actions, earthquakes, civil disturbances and chemical accidents.

While we have established that fire definitely presents a significant social and economic threat, the question to be addressed here is whether an expert system can aid fire emergency managers in more effectively utilizing available resources to mitigate the consequences of such fires and to diminish the likelihood of a conflagration all together. The challenge is to help fire managers get ahead of the fire, NOT chase it, by anticipating the condition of the fire over time and analyzing the effectiveness of alternative firefighting tactics.

To build an expert system, it must be possible to adequately describe fire behavior and the effectiveness of fire suppression efforts. While there are a number of influencing factors that effect fire behavior, fire is a reasonably predictable phenomenon. As is documented in research reports on the theory of fire and in experimental tests, there is a significant knowledge base and modeling experience to suggest that under certain known conditions, fire spread and growth is predictable to the degree of precision required by emergency managers although there remain some yet unexplained variation [2].

For instance, there have been apparent "discrepancies" in the damage due to fire following earthquakes. While some significant earthquakes such as the 1906 San Francisco quake generated great number of fires, other major tremors such as the 1971 San Fernando quake resulted in minimal damage. At this time, it is unclear whether these discrepancies are just statistical variation, differences in firefighting tactics or underlying differences in the built environment in the two cities.

Exhibit 1

Selected US Configurations (Source - NFPA and others)*[9]

Date	City	Buildings Destroyed	Loss (millions 1984 \$)*
Feb 7 1904	Baltimore	80 city blocks	500
Oct 23 1978	Malibu	230	60
Jul 26 1978	Santa Barbara	239	55
Oct 23 1947	Maine	1200	120
Jul 31 1979	Houston	26	25
Mar 21 1916	Paris TX	1440	80
Sep 17 1923	Berkeley CA	640	24
Oct 20 1944	Cleveland	89	30
Oct 23 1935	Los Angeles	222	20
May 15 1933	Auburn NY	250	12

* very approximate, intended for comparative purposes only

Although fire managers are frequently called upon to perform in emergency situations, their response efforts are largely confined to a single working fire. While an expert system can be of benefit in such situations, especially when the fire involves a complex building design or requires occupant evacuation, the value of an expert system increases as the number and magnitude of working fires increases such as may occur following an earthquake or in the case of forest and grasslands fires [8]. Here, the fire manager must make decisions regarding a relatively large geographic area and resource allocation with limited, and possibly, inaccurate information. The following description of the fire-related aspects of the 1983 Coalinga quake are included here to illustrate the emergency environment.

The following material is from a chapter in the Earthquake Engineering Research Institute report on the Coalinga Earthquake entitled "Fire-Related Aspects of the Coalinga Earthquake" [10].

"This chapter discusses the fires and related aspects of the Coalinga Earthquake of May 2, 1983. First reports from Coalinga indicated that fires caused by the earthquake were of conflagration proportions. Television news bulletins in San Francisco the evening of May 2 reported fires with "flames 40 and 50 feet in the air," a sight resembling "Tokyo in the great fire bombing." The San Francisco Chronicle of May 3 spoke of "...nearby oil pipelines ablaze,... buildings burst into flame." These reports exaggerated greatly the fire-related events of May 2, 1983. Nevertheless, events of that day contain valuable lessons with respect to the post-earthquake fire problem, which may be of importance in a similar or larger earthquake in a large metropolitan region. This chapter presents a narrative of the Coalinga earthquake, followed by a discussion of important observations and issues.

Fire-Related Events of May 2, 1983

The Coalinga Fire Department (CFD) consists of 30 personnel (11 paid, 19 volunteers). Equipment consists of two 1250 gpm and one 1000 gpm pumpers, one hose wagon and two ambulances, (the only ambulances in Coalinga, next nearest ambulances being Avenal which is 17 miles distant in the next county). Also located in Coalinga is a California Division of Forestry (CDF) forest battalion, equipped for brush fires. Located about 5 miles away is the Westside Fire Station. Following the earthquake, there were about 15 grass fires in the surrounding countryside, caused by downed or arcing power lines. These grass fires are not reported as having caused extensive damage. There were no reported oil-related fires. At the time the earthquake occurred (4:42 pm, PDT), there were four personnel in the fire station (Capt., Lt., firefighter and ambulance attendant). When the shaking began, they exited the building and saw collapses of nearby buildings. After the shaking had stopped, they moved the CFD equipment out of the fire station and issued a radio All Call (ie, all CFD personnel report to the fire station). The Captain (John Donelan) proceeded in his radio-equipped car to perform a reconnaissance of the Central Business District (CBD), also known as the Coalinga Plaza and located two blocks away. Large

amounts of dust had been raised, which obscured the CBD from the view of the fire station, so that the extent of the damage was unrealized. Arriving at the Plaza, Donelan radioed to the station, declaring a disaster and requesting notification of OES and mutual aid fire companies. At this time he was located at 5th and Elm, surrounded by brick buildings, when an after-shock struck (3 min., 46 sec, after the main shock, times taken from the CFD radio log). He then relocated to 5th and Durian, a safer location, and continued setting up a command structure and notification of outside authorities (relayed by radio). By this time, many brick buildings had collapsed, with consequent brick debris as much as 20' into the street on either side, measured from the building face. However, Coalinga building-to-building street width is 80' (wider than average) so that streets were still passable in the center.

About 8 minutes after the main shock, as dust was clearing, smoke was observed from the CBD and CFD responded with two engines, with however only one man on each. Hydrants and hose layouts were assigned, downed power lines identified and fire fighting commenced. The fire was in 260 Coalinga Plaza, the Coalinga Inn, a two story brick building which had partially collapsed and on arrival resembled a "bonfire" rather than a building fire. Due to the shortage of CFD personnel, citizen volunteers manned hose lines with a professional firefighter at the nozzle. Immediately after the main shock, Calif. Div. Forestry (CDF) personnel and trucks (which are equipped for bush fires) were dispatched for triage within Coalinga. CDF personnel decided to respond to Coalinga rather than to the above-mentioned grass fires based on likelihood of injuries in Coalinga and relatively low fire hazard in the grass in early May. Their officers assisted in mutual aid coordination. About 20 minutes after the main shock Westside fire units arrived and assisted CFD (Westside Fire District has three fire stations within twenty miles of Coalinga). The fire, although not especially large, proved extremely difficult to extinguish. At the time of this writing, CFD judgement is that the fire was fueled by high-proof liquor from broken bottles in a storeroom on the first floor, which also contained a water heater. This is based on interviews of the owner and review of the debris. The persistence of the fire is attributed to its being a flammable liquid fire, to hose streams being deflected by the debris, and to the "open-ness" of the partially collapsed structure. The Coalinga Inn is a relatively small building (20 feet by 90 feet) and efforts were also hampered by aftershocks and collapse of buildings. There was a standing order not to enter buildings for fire fighting purposes, due to the danger of collapse. After 4 hours, the fire had spread to the contiguous building (Petty's Jewelry Store), threatened the next building (Pressy's, housing most of Coalinga's insurance records), and had twice ignited the Webb building (both times extinguished). The Webb building is a two-story brick building separated from the Coalinga Inn by an alley (20 foot width). It had damaged but still intact walls during these events (the second floor walls were down the next day). Due to the stubbornness of the fire, it was finally decided, 5-1/2 hours after the main shock, to demolish Petty's Jewelry Store. That is, removal of the debris would constitute a fire break. This was initially attempted using a crane, which proved too slow, and finally accomplished using a large

front end loader and dump trucks.

During these 5-1/2 hours there also occurred:

- Chemical spills at the West Hills Community College and the Coalinga High School and a report of chemical spill at a dry cleaning establishment.
- Fire alarms (but no fire) at the Coalinga Convalescent Center and the Coalinga Hospital.
- Several attempts, finally successful, to verify that the town's gas had been shut off (Coalinga has a municipal gas system). Verification of this took 3-1/2 hours. Interviews with Coalinga residents indicate that many people on their own initiative turned their own and neighbors' gas off.
- There was an electrical outage caused by the main shock. By and large, residential electricity was not turned off house-by-house at personal initiative however, in contrast to the experience with gas.
- CFD and other vehicles sustained flat tires.
- Shortages of fuel for CFD vehicles, and shortages of drinking water at the fireground.
- A second fire, at 205 South Joaquin St. This was a residential kitchen fire, reported to CFD by telephone at about 8 pm. The cause is attributed to the following: shelf items were thrown by the earthquake around the kitchen, striking the electric stove controls and also landing on the heating elements. Electric power was later restored without notifying CFD. Damage was largely confined in the kitchen.

A third fire (510 South Joaquin St.) also occurred at about 8 pm. The cause was very similar to the fire at 205 South Joaquin but people were present and quickly extinguished the fire, with minimal damage. A fourth fire at 1080 Joaquin St., also a residential kitchen fire, occurred at about midnight. The cause was an electrical appliance (a can opener) being turned on by the flying debris, overheating and setting debris on fire. Damage included the entire kitchen and heat and smoke damage throughout the house. This house belonged to a police officer who was responding to the earthquake emergency at the time. It is significant that all these fires on Joaquin and South Joaquin Sts. had similar causes. Another potential ignition source was, six hours after the earthquake, down and arcing wires reported behind the Bank of America in the Coalinga Plaza."

This description provides an illustrative example of the need to integrate and interpret data from many sources and to make decisions based on erroneous information. The potential benefit of an expert system is also obvious.

II. Expert Systems Design

How can a computer system can help? In general, an expert system is not only to retrieve and integrate information but also to analyze and interpret it. An expert systems approach involves the definition of the possible decisions, the identification of influencing factors, the creation of a knowledge base and the application of one or more mechanisms for drawing inferences. For instance, decisions can involve evacuation, fire suppression tactics and fire containment. Fire loss is influenced by ignition frequency, conflagration potential and fire suppression resources. Finally, the creation of a knowledge base may combine results from laboratory test data, fire experience and fire models.

The following list presents the information that would assist fire managers at both the local and regional levels. At the regional level, an expert system could provide the following information:

1. Identify urban areas in the United States that are prone to conflagration under the prevailing weather conditions.
2. Evaluate the effectiveness of hardening buildings and the water supply in conflagration prone areas.
3. Evaluate tactics that break (interrupt) the growth and spread of fire when Forest Service personnel and/or the National Guard are available to aid local fire companies.
4. Review alternative tactical responses to firefighting under various resource allocation schemes.
5. Determine alternative mechanisms for either reducing average fire loss or minimizing the likelihood of a conflagration, such as the installation of sprinkler systems.

At the local level, an expert system could provide the following information:

1. Integrate data provided by a large number of widely dispersed sources.
2. Display confirmed ignitions and their magnitude and project the fire growth profile.
3. Display likely ignitions sources such as locations of known open flames and chemical mixing.
4. List the closest available fire fighting equipment and suppression agents for each confirmed ignition.
5. Display inaccessible areas.
6. Forecast (predict) possible fire spread scenarios.

7. Identify locations of likely ignitions that will be isolated by major roadway failure.

This information is essential for supporting the decision process at the fire scene. With respect to a specific emergency, there are the following questions regarding firefighting tactics:

1. What resources should be dispatched to each fire over time?
2. What is the proper fire fighting task--contain the fire spread, extinguish all burning, or perform rescue activities?

When considering these items, most involve either the integration of information from several different sources and/or the interpretation of information. As a result, it seems that an expert system, with components that may vary in the degree to which "artificial intelligence" is used, would be of benefit to the fire manager.

The critical requirements for developing an expert system concern the existence of a knowledge base and an inference mechanism. The knowledge base can be constructed from several sources including fire tests, historical data and fire models. Fire protection as an area of scientific study and engineering practice provides many of the theoretical relationships. There are major laboratories such as the National Bureau of Standards, Factory Mutual Research Corporation, Sandia Laboratory and Lawrence Livermore that have developed considerable understanding of the physics and chemistry of fire [1,7 and 11]. This paper will present several models that use a concept of fire development realms to describe the phases of fire behavior. The fire data along with the modeling abstraction of realms provide the elements for developing a knowledge base.

Drawing a conclusion about the effectiveness of alternative courses of action is very complex and requires an assessment of both fire behavior as well as suppression activities. A reasonable model of fire emergency management might follow the same format as the "RUN CARD" as illustrated in Exhibit 2 that typically prescribes the response for five alarm levels. Each alarm level reflects the known hazards along with the risks associated with events not yet realized because of inaccurate or incomplete information. The following factors are among the most influential:

1. Expected fire size,
2. Number of persons endangered,
3. Potential property loss,
4. Environmental conditions, and
5. Conflagration potential.

For instance, fire departments may make major commitments of manpower to a grassland fire because of the high possibility of unbridled spread even though the fire might be thought to be very small. In a specific situation, the tactical decisions are complicated by the following

Exhibit 2

--- RESPONSE CARD ---

Cooper River Drive

1121 to 1184

Engineers		Trucks	Chief	Special	Move-Up
6	4	T4	3-2-4-1 2-4-1		////////////////////
10	103	T10	Duty Chief Dep. C		15-10 7-6
7	3	T3	Chief		2-3 18-6
18	13	T13			17-6 27-13
17	1	T1			5-1 75-1 19-5

uncertainties:

1. Number, size and locations of fires known at this time,
2. Accessibility to the identified fire locations,
3. Adequacy of fire suppression agents (water, chemicals, etc.), and
4. Sufficient fire personnel for the assigned tasks.

The correct assessment of these uncertainties is critical to the choice of firefighting tactics, which in turn requires the rapid integration and interpretation of data for disparate sources.

Now the question concerns the most appropriate design of an inference mechanism. It appears that fire, like many other real problems, requires a hybrid of the major approaches to expert systems design. There are a number of questions that fit the "IF-THEN" model such as "If there are persons trapped, then initiate evacuation" or "If the chemical is chlorine, then do not use water." However, the identification of the appropriate evacuation routes must be determined by a mathematical model that considers the status of the street network at present and in light of the potential fire spread. Also, decisions based on the rate and size of the fire in the next ten minutes are best approached using a simulation model.

III. Fire Modeling Elements of an Expert System

There are computer tools presently available to assist fire managers. For illustrative purposes, three different models are briefly discussed. The first is FIREMOD which simulates grassland fire spread. The second is EMBER which models fire spread and occupant movement within a building. Finally, FILO is an urban fire spread model. In many ways the models are similar in the mechanisms used to describe fire growth and spread. None of these models is designed to tell a fire manager how to fight a fire.

FIREMOD - Wildlands Fire Spread

FIREMOD is a computer program designed by the U.S. Forest Service to project wildland fire spread potential and to assess the probable success of initial attack forces. The projection of forward and lateral fire spread can be obtained for one-hour increments without recognizing any suppression activity. FIREMOD is not designed to tell a fire manager how to fight a fire but the direction and speed of spread. FIREMOD uses the following information:

1. Fuel type (grass, chamise or brush)
2. Slope
3. Age of fuel

4. Fine fuel moisture

5. Wind speed

Based on these input data, the elliptical spread model displays the following information:

1. Expected forward rate of spread.
2. Expected ellipse dimensions.
3. Flame length.
4. Heat energy output of the fireline.

The effectiveness of alternative firefighting tactics are displayed in Exhibit 3.

Because of the uncertainty in the number and location of spot fires, some departments faced with grassland fires will commit a major part of the firefighting resources to a suspected fire. In this way, the precise deployment of units may not be critical since there will be sufficient manpower to survey the entire area. Furthermore, there is likely to be sufficient firefighting power to extinguish quickly the incipient fires. Time is of utmost importance for these types of fires since there are few natural barriers.

Exhibit 3: FIRE SUPPRESSION INTERPRETATIONS

- I. Fire Line Intensity: 100 BTU/SEC/FT
Flame Lengths: 4 Feet
Firefighting Tactics:

Fire can generally be attacked at the head using hand tools.

- II. Fire Line Intensity: 100-500 BTU/SEC/FT
Flame Lengths: 4-8 Feet
Firefighting Tactics:

Fires are too intense for direct attack on the head
by persons using hand tools.
Hand lines cannot be relied on to hold fire.
Engines, dozers and retardant aircraft can be effective.

- III. Fire Line Intensity: 500-1000 BTU/SEC/FT
Flame Lengths: 8-11 Feet
Firefighting Tactics:

Fire may present serious control problems, i.e. crowning
and spotting.
Control efforts at the fire head will probably be ineffective.

- IV. Fire Line Intensity: 1000 BTU/SEC/FT
Flame Lengths: 11 Feet
Firefighting Tactics:

Crowning, spotting and major fire runs are probable.
Control efforts at the head of the fire are ineffective.

EMBER - Building Firesafety Model

EMBER was designed to analyze the hazards of fire in a structure. Here the fire manager is often faced with the dilemma either to extinguish the fire thereby both protecting the building occupants and saving the structure itself or to assist occupants in reaching a location of safety and then to extinguish or confine fire spread. Unlike the wildland fire situation, the building structure provides barriers to fire spread that enables the fire manager to consider alternatives that require less resources.

A preliminary version of this computer model was initially developed with support from the Department of Housing and Urban Development. This model can be used to determine the available escape routes, the expected time for evacuation and the anticipated extent of fire spread and toxic levels of the combustion gases.

EMBER is a computer model that simulates the movement of smoke and occupants as though an experiment were being conducted [3]. Fire development within a building is described in terms of "realms." A fire realm describes a phase in the development of a fire with regard to its heat release and flame size as illustrated in Exhibit 4. Exhibit 5 presents data depicting the average time in several realms obtained from full scale fire tests for two dwelling types--mobile home and single family design. Because these data were drawn from fire tests for which it was desirable to save the structure, all fires were extinguished at the last moment. As a result, data were not available for the higher realms. Based on a description of the building, fire profile and occupant locations, the model provides a continuous record of temperature and carbon monoxide levels in each building space as well as the occupant locations and their exposure to carbon monoxide.

The purpose of this model is to describe the potential for normal and handicapped persons to escape or be rescued in the event of a fire or similar emergency. The modeling approach involves the construction of an egress and airflow network for a building and computer simulation. The "egress network" is a graphic representation which describes the connections among adjacent spaces. The "airflow" network adds channels for airflow to the egress network. Based on this combined network, a computer analysis is used to evaluate the potential for occupants to escape.

The evacuation analysis of two prototypical transit stations is used to illustrate the application of EMBER. The focus is on the identification of those elements of station design that are essential for rapid and reliable evacuation. The analysis involves the following four steps:

- Step 1: Obtain Station Designs
- Step 2: Construct Egress and Airflow Networks
- Step 3: Distribute Occupant Load
- Step 4: Run EMBER

Exhibit 6 illustrates the longitudinal section for two prototypical transit stations that are approximately 40 feet below the surface with platforms 450 feet long and 30 feet wide. The first station is a double

Exhibit 4: REALM III


Critical Events

Indicators	Value
=====	
Flame Height	.3 m
Flux to Fuel	
from flame	20 kW/m ²
from ceiling	1 kW/m ²
Power Level	5 kW
Remote Flux	0
Gas Temperature	15 ⁰ k rise

Exhibit 5: REALM Profile

Duration (seconds)				
	II	III	IV	V
=====				
Mobile Home	200	251	105	?
Single Family	262	1041	?	?

DOUBLE LOADED MEZZANINE PLAN



TYPICAL LONGITUDINAL SECTION

end-loaded design such that evacuation could be along two stair towers at opposing ends of the platform and on the two combination stair and escalator paths from the platform through the mezzanine to the surface. The second is a single end-loaded station. In this design evacuation is through the single stair tower at the end of the platform or the tandem escalator and stair arrangement at the other end.

Exhibit 7 and 8 illustrate a portion of the network for the station platform. Here, the use of nodes to represent a train car as well as the adjacent platform area is illustrated.

To evaluate the dynamic aspects of occupant evacuation, EMBER is run with three different occupant scenarios: (1) a light occupant load of 640 persons distributed on two trains and the platform; (2) a heavy load of 3142 persons distributed among a single train, the platform and mezzanine; and (3) a crush load of 4188 persons divided among two trains, the platform and mezzanine. To illustrate the potential impact of disruptions to the circulation system, the stair tower and two interior paths are blocked. Table 1 illustrates selected results from these EMBER simulation experiments.


The following observations concern the escape potential for these two station designs based on the simulation results:

1. The major bottleneck in both designs is the stair tower. While it is required to comply with the fire code provisions regarding remoteness of exits, the stair is actually detrimental at the present dimensions. To provide both quality and reliability, the stair and door widths should be increased.
2. Stair tower(s) alone are insufficient to provide for the rapid evacuation of the platform should the mezzanine be blocked which diminishes the value of these stairs as a reliable, separate means for evacuation.
3. All escape routes are well used making it very difficult for transit personnel and firefighters to travel in the opposite direction to reach the fire or control equipment located inside the station.
4. Stair towers will not provide a smoke-free environment in the event of a fire since the platform doors remain open during most of the evacuation.
5. The percentage of persons that are closer to a stair tower under crush load conditions for the double loaded station plan is 37%. This proximity to such a large proportion of the occupants explains, in part, why the stairs play such a vital role in the rapid evacuation process.
6. There is a potential for misdirected or dead-end travel by having doors opening into the ancillary rooms at the end of the platform located near the platform entrances to the emergency stair towers.

SINGLE END LOADED PLAN

Node Key

- A Ancillary
- E Escallator
- M Mezzanine
- P Platform
- S Stair
- T Train
- X Exit



DOUBLE LOADED MEZZANINE PLAN




Table 1: Simulated Evacuation Times

Station Design	Evacuation Times (seconds)	
	Light Load	Crush Load
=====		
Single End Loaded	425	1414
Double End Loaded	253	1341
Single End Loaded with Blocked Stairway	346	
Double End Loaded with Blocked Stairway	298	
Single End Loaded with Wider Stairway	313	1240
Double End Loaded with Wider Stairway	226	1086

Table 1 also presents the impact of various disruptions to the circulation system such as when the stair tower is blocked and alternatively, increased in width. When the single stair tower is blocked, evacuation time is increased approximately 19% for the double end-loaded station but decreased for the single end-loaded station. The decrease in evacuation time that resulted from the blockage of an exit illustrates the complexity of the evacuation dynamics. When the stair tower is blocked, occupants do not commit themselves to using the nearby stair which rapidly becomes congested. Instead, occupants follow a path through the mezzanine despite the fact that it seems at first to be a longer route, but because of the larger capacity, actually results in a more rapid evacuation.

This type of model can be used to analyze different strategies for gaining access to the fire without seriously disrupting the occupant evacuation and for identifying critical locations that should be monitored so that additional congestion does not occur.

FILO - Urban Fire Loss Simulation

This is a discrete event simulation based on a description of the built environment and climatic conditions. This computer model was initially developed with support from the National Science Foundation. Initial fires are located according to a probability distribution that is determined by the magnitude of the causal condition such as the intensity of an earthquake. The FILO simulation then describes fire development within a building and its spread from a building to its neighbors.

For each building, the following data elements are required:

1. The Building Number
2. Building Type and Location
3. Realm Profile for each Building Type
4. Location of Active Fire Companies


Exhibit 9 illustrates the type of information that can be assembled for a typical building type.

In addition, two global quantities will be specified, namely the wind velocity and humidity. These quantities have a dramatic influence on impingement possibilities and radiation effects.

Given the above data, the program will:

1. Calculate the distances between adjacent buildings.
2. Compute the expected times to enter the next fire realm.
3. Compute the time of spread to an adjacent building.

Exhibit 9: TYPICAL SAN DIEGO SINGLE FAMILY RESIDENCE



ROOMS :

B Bath
K Kitchen
L Living Rm
D Dining Rm
BR Bedroom
E Entry Hall
G Garage
U Utility Rm

FIRE SOURCES :

1. Gas-fired water heater (unanchored)
2. Air conditioning unit
3. Shelf storage of possibly flammable liquids
4. Automobiles
5. Electric space heaters
6. Wood burning fireplace
7. Electric clothes dryer
8. Gas range
9. Outdoor gas grill

CONSTRUCTION :

Windows: Double glazing in anodized aluminum sash and frame

Insulation: 4" fiberglass batt in walls 8" in ceiling

Floor : 6" concr. slab 8" slab in garage

Fl.Fin: carpet or vinyl asbestos tile

Ceiling: 5/8" gypsum board

Roof: Plywood on wood trusses or wd.rafters

Interior Walls: 5/8" gypsum over wood studs

At Utility Rm: 2 layers of 5/8" gyp. over studs on kitchen side; and 1 layer on utility room side

At Garage: 8" concr. masonry wall between garage and entry & utility rm fin at entry with 5/8" gyp. on furring

Exterior Walls:

Wood studs with 5/8" gyp. fin on interior. Exterior finished with 3/4" stucco over plywood sheathing

4. Assign the closest available fire companies to each fire.

Taken from a similar model developed by Scawthorn [9], Exhibit 10 illustrates a CRT screen display and Table 2 illustrates the tabular report on the fire spread.

IV. Implementation of an Expert System Approach

These three models can be used to respond to a wide range of hypothetical and actual situations. However, each provides results that are dependent upon a number of specific conditions. Since there are an infinite number of input values for each of the pertinent parameters, there is potentially an infinite number of different outcomes. However, the analysis procedure remains the same, and hence, amenable to an expert systems approach.

Thus, an expert system for fire emergency management will rely heavily on mathematical modeling because of the dynamic and spatial aspects of fire growth and resource allocation. It will also embody both the concepts and elements of a Decision Support System and Artificial Intelligence. This structure also suggests that the expert system shell should enable rule-based, frame-based and object-oriented representation as well as forward and backward logical chaining. For instance, backward chaining would be used to identify the fire site for several observations of smoke while forward chaining would be used to describe the expected resource assignment for a confirmed ignition.

Furthermore, the inference mechanism provided by existing software such as the Automated Reasoning Tool (ART) from the Inference Corporation is advertised as being as to process "viewpoints". Viewpoints allow ART to consider potentially conflicting alternatives such as the conflict between extinguishment (Defend in Place) and evacuation (Move the Exposed) strategies toward lifesafety. The capabilities of this particular expert system shell also enables the system to explore many aspects of one problem as well as to evaluate a scenario over time.

This application focuses on the acquisition, display and interpretation of data on the potential number of persons effected, the location and magnitude of the fire and the availability of firefighting resources. Fortunately, much of the the required data and technology are readily available to support this application of an expert system.

1. DIME File. This is a data base originally designed by the Census Bureau that describes the road system for most urban areas in the United States. The DIME file provides the means of calculating the possible paths between one location and other. Used in conjunction with information on road segments that are impassable provides a mechanism for calculating the difficulty to reach a particular area if a fire does occur.
2. Census Demographic Data provides information on the number of persons, number of housing units and average value.

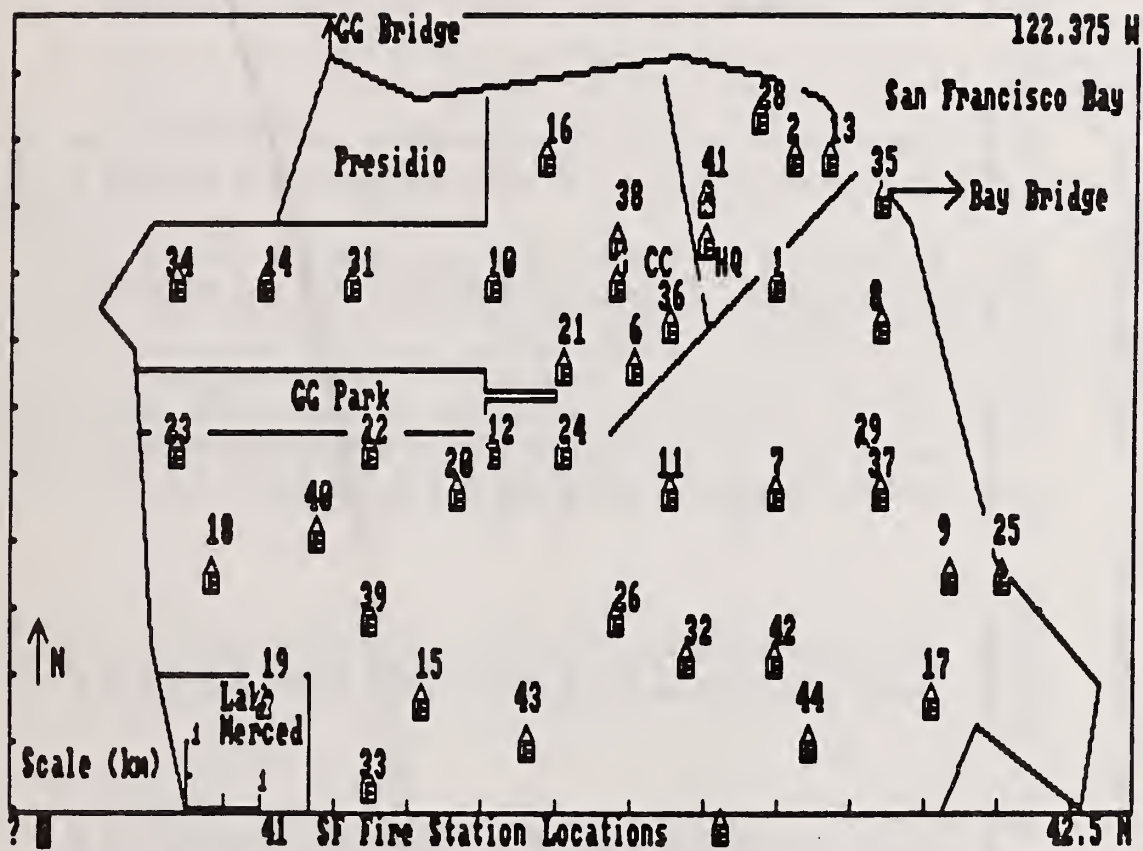


Exhibit 10: San Francisco Fire Department Fire Station Locations

Table 2 Magnitude 8.3 San Andreas Event Scenario,
San Francisco: Initial Fire Department Response and Demand

Wind velocity 10 mph from azimuth 270 INITIAL RESPONSE									
Fire	T(arrival) (mins)	Fire Vel. (ft/min)	A(Burnt) (acres)	No. Burnt Bldgs	Loss (\$mil)	FFlow (Mgpm)	No.Engs. Reqd.	No. Pers. Reqd.	
1	10.5	9.2	0.2	2.2	0.2	1.3	2	11	
2	6.2	7.6	0.1	1.6	0.1	1.0	1	8	
3	8.1	8.4	0.1	2.8	0.2	1.7	2	14	
4	7.4	8.1	0.1	1.1	0.1	0.6	1	5	
5	8.5	8.5	0.1	1.3	0.1	0.8	1	7	
6	19.9	11.4	0.8	18.7	1.4	11.1	15	93	
7	21.8	11.7	1.0	22.9	1.7	13.6	18	113	
8	6.2	7.6	0.1	1.6	0.1	1.0	1	8	
9	15.2	10.6	0.4	3.8	0.3	2.3	3	19	
10	8.1	8.4	0.1	2.8	0.2	1.7	2	14	
11	11.6	9.6	0.3	2.4	0.2	1.4	2	12	
12	9.7	9.0	0.2	10.0	0.7	5.9	8	49	
13	8.1	8.5	0.1	6.4	0.5	3.8	5	32	
14	15.6	10.4	0.6	26.0	2.0	15.5	21	129	
15	7.4	8.2	0.1	5.0	0.4	2.9	4	25	
16	10.5	9.2	0.2	2.2	0.2	1.3	2	11	
17	10.5	9.2	0.2	10.0	0.7	5.9	8	49	
18	15.2	10.3	0.5	22.8	1.7	13.5	18	113	
19	8.5	4.9	0.1	2.5	0.2	1.5	2	12	
20	12.4	9.7	0.3	14.4	1.1	8.6	11	71	
21	5.0	6.2	0.1	5.6	0.4	3.3	4	28	
22	6.2	4.7	0.0	1.5	0.1	0.9	1	8	
23	9.3	8.8	0.1	1.3	0.1	0.8	1	7	
24	14.7	10.7	0.3	1.6	0.1	1.0	1	8	
25	9.3	8.7	0.1	0.6	0.0	0.3	0	3	
26	9.3	8.8	0.1	1.3	0.1	0.8	1	7	
27	11.6	9.7	0.2	1.0	0.1	0.6	1	5	
							137	859	

3. Weather conditions regarding wind direction and speed, temperature and humidity are readily available.
4. Radio and remote phone equipment is available to enable communications among all personnel.
5. Affordable large format CRT terminals and large screen projection systems enable the implementation of multi-dimensional graphical displays.
6. Fire realms and "alarm level" abstractions provide a manageable mechanism for describing fire behavior and suppression response requirements.

These elements provide the mechanics required for the practical aspects of implementing an expert system.

Since emergency fire management is often carried out under hectic conditions with limited information and scarce resources, mistakes in early tactics and resource deployment can be costly. The three models presented here illustrate techniques that can be employed by personnel lacking formal fire engineering training to participate in emergency management in order to assess the risks and allocation resources.

V. References

- [1] Alpert R.L., A.T. Modak and J.S. Newman, "The Third Full-Scale Bedroom Test of the Home Fire Project," Volume 1, Factory Mutual Research Corporation, Norwood, Massachusetts, 1975.
- [2] Berlin, G.N., "Managing the Variability of Fire Behavior", Fire Technology, Volume 16, No. 4, 1980.
- [3] Berlin, G.N., "A System for Describing the Expected Hazards of Building Fires," Fire Safety Journal, Volume 2, 1980.
- [4] Berlin, G.N. and William T. Hathaway, "Transit Station Evacuation," National Bureau of Standards, Gaithersburg, MD, 1985.
- [5] "Fire Experience and Exposure in Fixed-Guideway Transit Systems," Report written by Schirmer Engineering Corporation and Rolf Jensen & Associates, Inc. and prepared for the American Iron and Steel Institute, December, 1980.
- [6] Hathaway, William T., "Fire Safety in Transit Systems," Paper presented at the Fifth International System Safety Conference, Denver, Colorado, July, 1981.
- [7] Quintiere, James, "Some Observations on Building Corridor Fires," Fifteenth International Symposium on Combustion, The Combustion Institute, August 1974.
- [8] Scawthorn, C., Y. Yamada, H. Iemura, "A Model for Urban Post-Earthquake Fire Hazard," DISASTERS, Volume 5, No. 2, London, England, 1981.
- [9] Scawthorn, C., "Post-Earthquake Fire Considerations in Urban Earthquake Policy and Planning," U.S.-Japan Workshop on Urban Hazards Reduction, July, 1984.
- [10] Scawthorn, C., and John Donelan, Research Institute report on the Coalinga Earthquake entitled "Fire-Related Aspects of the Coalinga Earthquake," Earthquake Engineering Research Institute, Berkeley, 1984.
- [11] Tewarson, A. "Final Report - Experimental Evaluation of Flammability Parameters of Polyment Materials," Factory Mutual Research Corporation, Norwood, Massachusetts, 1979.

Acknowledgement

I wish to express my appreciation to Dr. Charles Scawthorn of Dames & Moore for his many contributions to this paper and to John Gerard of the National Fire Protection Association.

The VENTILATION EXPERT for Underground Coal Mines

by

Jeffery L. Kohler
Assistant Professor of Mining Engineering
The Pennsylvania State University
Department of Mineral Engineering
University Park, PA 16802

ABSTRACT

The underground coal mine environment does not present health or safety risks to workers when it is maintained to specific standards. Mining activity, as well as unforeseen events, can result in sudden and dangerous levels of explosive gas and dust. A mine may have hundreds of miles of passages spread over ten or more square miles. Control devices for maintaining the environment are distributed throughout the infrastructure of the mine. Many of the decisions which have to be made to maintain a safe environment are based on judgment, rather than documented facts or empirical procedures. All of these characteristics of the system compound to make it difficult to manage ventilation under routine conditions. When a fire or explosion occurs the management problem escalates dramatically.

The development of a prototype expert system to manage the operation of coal mine ventilation systems is described in this paper. The characteristics of the mine, as they affect the expert system design are presented. The knowledge engineering aspects of this system, as well as implementation issues are examined. The paper is concluded with some observations on the requirements of a future version of the VENTILATION EXPERT.

INTRODUCTION

The Pennsylvania State University, Department of Mineral Engineering was funded by the Bureau of Mines to perform a conceptual investigation of a management information system for coal mine ventilation systems (1). Initial work under this contract suggested that an expert system implementation would be most appropriate. Although the actual construction of an expert system was not called for under the contract scope-of-work, it was decided that a small prototype system should be built and evaluated. This paper is a summary of the ongoing effort to build and evaluate expert system structures for the coal mine domain. The effort described here is also one in which the domain experts performed the development work, rather than knowledge engineers.

The design objectives of the management information system are presented in the next section of this paper; then the coal mine domain is briefly reviewed as it relates to the expert system. This is followed by a discussion of the selection criteria applied to the

selection of the expert system tool. The general aspects of the selected architecture are then reviewed as they relate to the coal mine problem. The paper is concluded with a summary of the experiences, including the strengths and weaknesses of the selected control structure; planned changes for a future implementation are also given.

DESIGN OBJECTIVES

The objective of building a prototype was to achieve a better understanding of what a coal mine expert system should be capable of performing; how it should be structured; and the software and hardware constraints that might be applicable. Since the prototype system would only be used by the researchers and not mine personnel, some additional flexibility existed in areas such as "user friendliness," throughput rate, and so forth. Notwithstanding, the objectives of the future system were carefully defined to allow for a more realistic formulation of the prototype system.

The objectives of the system generally correspond to removal of deficiencies in the operational methodology of many existing mine ventilation systems, and more specifically, perceived deficiencies as computerized monitoring systems become more commonplace in mining. The types of information-management problems which will become more commonplace as monitoring systems become integrated into mine operation can already be observed in the approximately 56 mines where monitoring systems have been installed.

The most notable problem is the large amount of data generated by the monitoring system and other sources, and the relatively small amount of information contained in that data. The sheer bulk of data discourages personnel from using it, thereby negating many of the benefits of monitoring; a more serious consequence is the possibility that important information within the data will not be noticed. Thus an important objective of the proposed management information system is to sift and sort the data for information useful to specific tasks and individuals. Secondly, this information should be automatically routed to the appropriate personnel.

An examination of decision making, with regards to mine ventilation, reveals other needs which should be incorporated into the design objectives for the system. The first aspect of decision making is the time frame. Some decisions must be made in a matter of minutes or even seconds, as during a fire or after an explosion in the mine. Others can be made over weeks or months; for example long term improvement of ventilation system performance. Still others fall somewhere in between these two extremes. However it is the extremes which illustrate the need for a system which can provide not just information in short time, but also advice based on an intelligent use of the information. In the one case there is so little time to make the decision that all resources must be brought to bear, while in the other extreme so much time has elapsed that information may have become scattered or lost. Thus, another objective of the system is to utilize the available information to provide advice to management.

The resources available to support the decision-making process include mathematical relationships such as the fan laws, computer simulations, statutory requirements, and the judgment and experience of the decision maker. An obvious objective of the system is to embody these same resources. Embodying the system with the capability to utilize heuristical information is particularly important since many mining decisions have to be based on heuristics. Furthermore, the people with the experience are not always available at a local level to participate in the decision-making process. It would therefore be very desirable if this experience could be extracted from the expert, encoded, and then made available to everyone through a knowledge-based expert system.

Given the minimal time to make certain decisions, and the severity of the consequences if they are not made promptly and correctly, it would appear desirable to allow the system to assert control as appropriate and necessary. For example if the system's advice is to change fan speed, and mine personnel have not responded, then it may be feasible to allow the system to change the speed through a control loop. This would be illegal under current law and raises other issues, but is a germane issue at this stage of the research.

The aforementioned objectives can be summarized as a set of functional goals; the system should have the capability to:

- manage the information flow to personnel,
- advise management on problems and solutions, and
- make decisions and assert control.

The nature of the information, the functional goals, and the need to allow for the addition of user-site-specific knowledge in the system suggest that a knowledge-based expert system implementation would be appropriate. Figure 1 illustrates the system concept. The actual implementation of the knowledge-based expert system (ES) is dependent upon the domain. A precise definition of the domain is not only necessary to structure an ES, but is useful in forcing a more precise understanding of the domain's interaction with the ES. The coal mine domain is briefly reviewed in the next section.

COAL MINE DOMAIN

An underground coal mine can be described in terms of several attributes; many of which are shown in Figure 2. These are discussed briefly in this section, with an emphasis on those which have a significant relationship to the ventilation system.

The infrastructure of a mine consists of a complex network of openings which have been excavated to facilitate access to the deposit, provide routes for men, materials, and ore, and to provide airways. The excavated openings are known as entries or crosscuts, depending on their orientation and function. Solid pillars of coal are left unmined to support the roof. Typically the width of an entry or crosscut is 16-20 ft., with entries being on 40-100 ft. centers, and crosscuts on

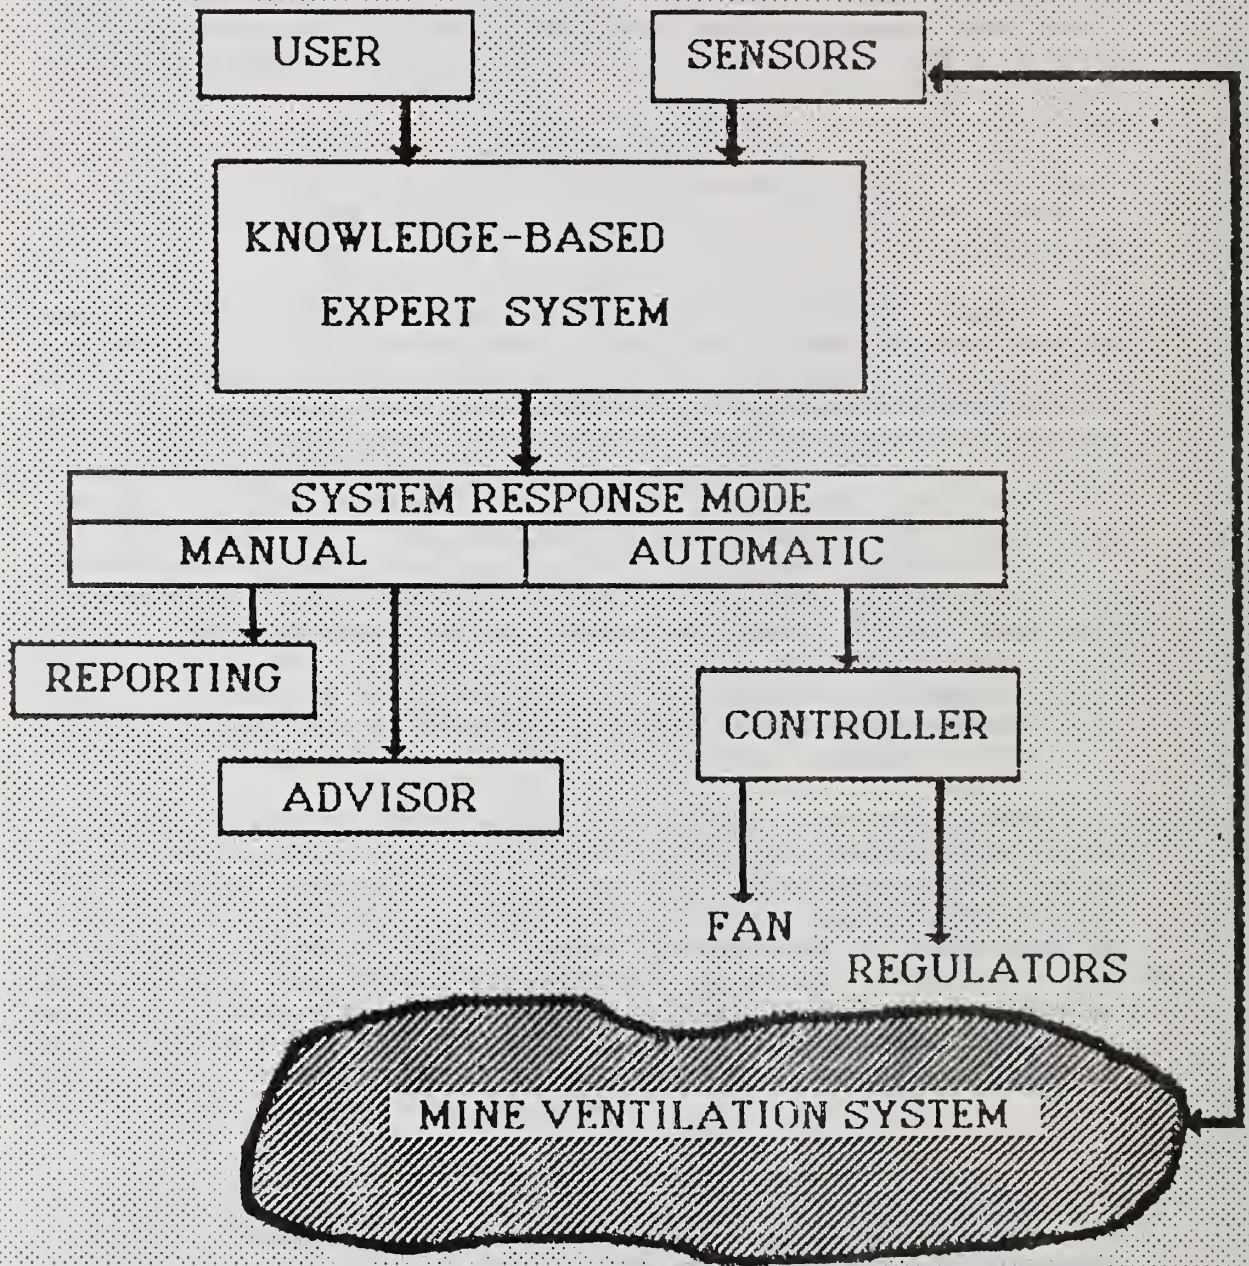


Figure 1. Ventilation Expert - System Concept.

60-150 ft. centers, depending on the pillar requirements. The height of an opening is usually the same as the thickness of the coal seam, and ranges from an average of 32 in. to 72 in., and even greater in the thick seams found in western U.S.


The entries are given specific names based on their function as airways and production areas, or their use for materials handling. For example if an entry is a major path to and from many parts of the mine it may be called a main entry. If that particular entry also carries intake air (fresh air) to the working place it would be known as a main intake entry. An example of the entry types is shown in Figure 3.

The deposit, an attribute of the mine, has its own attributes. Two of interest here are methane content of the coal and the structural integrity of the rock immediately above and below the coal seam. The methane content will determine the amount of methane liberated when the deposit is mined, and will affect the amount of gas which "bleeds" into the mine atmosphere. The structural characteristics are important insofar as the probability of a roof fall or floor heave is concerned; either a fall or heave will result in a blockage of the aircourse. The first attribute will tend to be similar throughout the mine, while the second may tend to exhibit more local variation.

The attributes of materials handling, production, power, and drainage constitute important subsystems within the mine. However, their import insofar as the expert system for mine ventilation is concerned, is minimal; accordingly they will not be discussed. The management attribute is important here because it defines the scope of the expert system and how it must interact with mine management to achieve the desired design goals.

The structure of mine management is determined by both statutory requirements and company organization. One example of the organization follows. At the local level a face boss is in charge of a discrete production unit, within the mine, known as a section. An assistant mine foreman has responsibility for a number of sections, and the mine foreman has responsibility for all the sections at the mine. The maintenance superintendent oversees all maintenance activities, and may have a foreman and other lower level managers reporting to him. The mine superintendent is in charge of the mine and has responsibility for all that occurs there. The mine engineer, safety engineer, and ventilation engineer report to the mine superintendent. Depending on the size of the company, the mine superintendents from several mines will report to a district manager, who reports to a vice-president for production.

The management attribute of the mine penetrates virtually every aspect of the expert system. The information available to the expert system, the interactive operation of the expert system, and so forth depends upon the management structure. A detailed account of this has been prepared, and is documented in reference 2. For the purposes of this paper, however, it is worth noting that the complex management structure serves to complicate the issue of information management as far as the monitoring system is concerned.



The ventilation attribute is central to the development of this expert system, and in its own right is one of the more important subsystems of the mine. The purpose of a coal mine ventilation system is to dilute, render harmless, and carry away all noxious and toxic gases and dusts. In some mines air conditioning is also performed to provide better temperature and humidity levels. The ventilation consists of entries carrying fresh air, i.e. intake entries, entries carrying spent air, i.e. return entries, a fan and other devices.

The importance of the ventilation system necessitates that it be examined in greater detail. Figure 4 illustrates the attributes of the ventilation system. For technical accuracy in the domain, it must be noted that face ventilation has not been included in this first generation of the Ventilation Expert.


The first three attributes of Figure 4 are control devices within the system. The fan is located on the surface and is the prime source of fresh air for the mine. A mine fan may have a capacity of 500,000 cfm at 10 inches water gauge for example. Large mines may have multiple fans. The operating point for quantity and pressure of the fan is adjusted to provide adequate air to the mine. Under current law this point cannot be changed by fan adjustment while workers are in the mine. For the purposes of this project it is assumed that the point can be changed at any time, to provide only needed air, thereby substantially reducing power costs, while always providing adequate ventilation.

Regulators are devices placed in certain aircourses to introduce a shock loss in the airflow, thereby increasing the resistance of that airway. These devices are used to control the ratios in which the air splits among a large number of entries. The technology has been demonstrated to control regulators remotely, and it is assumed that this capability will be available to the expert system.

Stoppings are walls placed in crosscuts between certain entries to separate different types of flows, for example intake from return. Their importance lies in their behavior and the attendant consequences when they begin to leak.

The environmental attribute of Figure 4 can be further quantified by illustrating its attributes, as was done in Figure 5. Each of these attributes can be sensed, and by monitoring them at appropriate points throughout the mine, the performance of the ventilation system can be precisely defined. Similarly, when problems arise, these parameters are useful for diagnostic purposes. Some examples of the significance attached to these measurable attributes follow.

Pressure and quantity measurements define flows throughout the mine. Besides comparing these to expected values, anomalies can be diagnosed by examining specific changes, or values, at different locations. Methane concentrations are significant with respect to preventing fires and explosions; additionally the concentrations in various locations are limited by statutory requirements. Carbon monoxide measurements are used to detect and to identify the source of combustion. In general the other attributes shown in Figure 5 would not



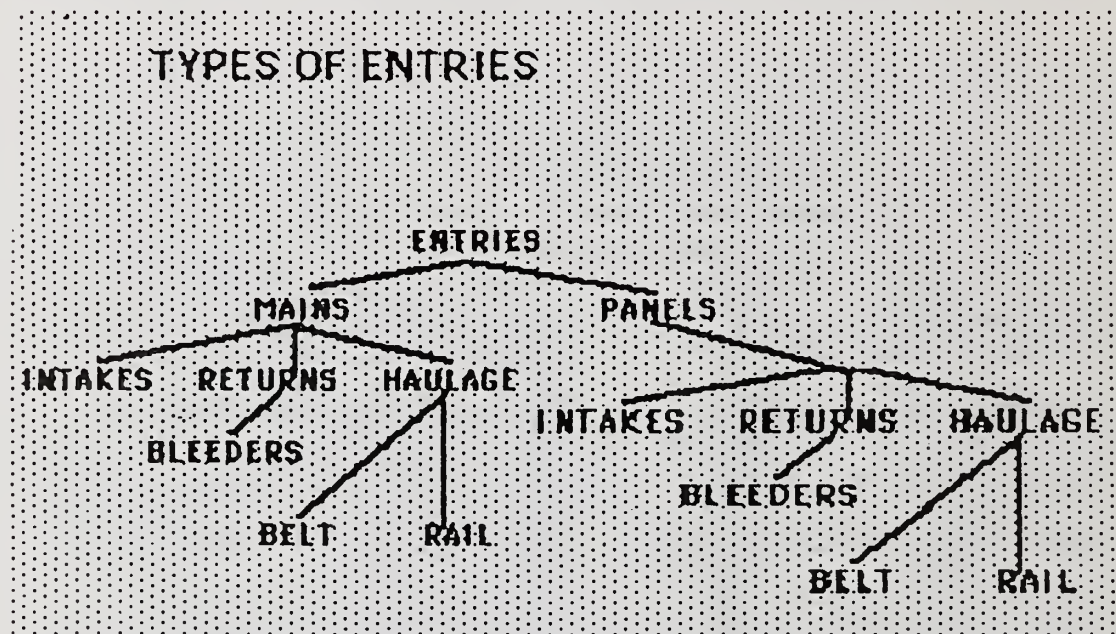


Figure 5. Environmental Attributes.

be measured, unless there was a site-specific reason. For instance if spontaneous combustion was a problem, then carbon dioxide might be sensed in gob areas for early warning of an incipient problem.

A properly designed and maintained ventilation system will provide a safe and healthy working environment without wasting energy. The mine is very dynamic, and as the infrastructure changes so does the ventilation system and the ventilation requirements. Ventilation requirements are also a function of production activities. In addition to changes in ventilation requirements for the aforementioned reasons, a number of other events can alter system performance. Examples include roof falls which may block aircourses, equipment moves which will obstruct aircourses, gas outbursts, and failure of labor to implement certain changes in the system such as regulator adjustment. Management response to these problems, as indicated earlier, may need to occur within minutes or over months, depending on the nature and severity of the problem.

For the purposes of this project, problems which the expert system will need to identify and react to have been divided into two classes: those occurring on a routine basis are grouped as normal and those which are catastrophic are grouped as emergency. "Normal" problems include inadequate air quantity, excess methane, and inadequate air velocity, among others. "Emergency" problems include fire, explosion, and gas inundation. It is important to note that the so-called "normal" problems are quite significant in that a failure to correct them can result in an explosion or fire. Additionally, there can be an adverse affect on production if they are not quickly corrected.

The requirements of the expert system and the domain in which it will be used have been summarized. With that as background, the criteria utilized for selecting an expert system tool will be described.

TOOL SELECTION CRITERIA

The primary objective of building the Ventilation Expert was to facilitate a more accurate specification of an expert system which would be available to mining companies at some future time. Since the objective was not to develop a system for general use, the criteria set for selecting a tool are somewhat different than might be the case otherwise. The applied criteria were:

- existing and demonstrated tool,
- readily available at low cost, and
- satisfy the intent of the research project.

The desire for an existing tool is obvious. The requirement for a demonstrated tool stemmed from the fact that we did not want to take time to debug a new tool; also, it was felt that a demonstrated tool would be more fully documented. The requirement on availability and cost was a reflection of the resources available for the effort. Finally, it was important that the selected tool be capable of doing the types of things which were perceived as important for the ventilation application. Satisfying this requirement was problematic, since an

implicit objective of building the prototype was to get a better idea of just what type of expert system architecture would be best suited to the ventilation problem.

The types of expert system tools available can be grouped into generic categories, for example: (3)

- interpretation, in which situation descriptions are inferred from sensor data;
- prediction, in which consequences of given situations are inferred;
- diagnosis, in which malfunctions are inferred from observables;
- monitoring, in which observations are compared to plan vulnerabilities;
- debugging, in which remedies are prescribed for malfunctions;
- repair, in which a plan to administer a prescribed remedy is executed;
- control, in which system behavior is interpreted, predicted, repaired, and monitored.

Utilizing these generic groups, the Ventilation Expert needs to first perform interpretation, then diagnosis, followed by debugging, and then prediction to verify the proposed remedy, prior to implementing the repair. The generic category of control essentially satisfies the Ventilation Expert's needs. Unfortunately, the control category probably has received the least amount of attention from those who develop expert system tools, and at the time this effort was begun a suitable expert system tool for control was unavailable. Consistent with the project goals and the other constraints on tool selection, a decision was made to use a production system known as mini-MYCIN. This particular system is quite similar to EMYCIN.

EXPERT SYSTEM IMPLEMENTATION

MYCIN is a rule-based expert system for providing diagnosis and recommended therapy for certain classes of infectious blood diseases. It is used in a consultation mode and has well-developed explanation capabilities. The solution method employed by MYCIN is backward chaining, from a diagnostic hypothesis to data, under the guidance of inferential rules. Knowledge is represented as productions, i.e. rules. Another important characteristic of MYCIN is its ability to reason under uncertainty. (3)

EMYCIN is the MYCIN expert system without the domain specific knowledge. Mini-MYCIN, developed at MIT, has the essential characteristics of EMYCIN, and was used to develop the Ventilation Expert. The development of an expert system using mini-MYCIN involves the development of the domain-specific knowledge base. The knowledge base consists of two major components: the structural component which defines the organization of the knowledge, and the problem solving component which provides the necessary knowledge to develop a solution method and make inferences or use other resources to solve problems. Specification of the structural component essentially means defining the

context types and parameters; specification of the problem solving component requires the development of rules which define the behavior of the domain, by capturing an expert's expertise about the domain.

Structural Component of the Knowledge Base

The ventilation expert uses three context types: mine, entry, and location. The relationship between context types is described by the static context tree which is shown diagrammatically by Figure 6. During execution of the Ventilation Expert, the context types are instantiated, as necessary, to form the dynamic context tree, which describes the relationship between specific instances of the context types for that particular execution (consultation).




Figure 6. Static Context Tree for the Ventilation Expert.

Parameters are characteristics of the context types. Representative parameters for the MINE context type are:

- Fan Characteristic Curve
- Fan Status
- Methane Liberation Rate
- Monitoring System
- Management
- Significant Problem
- Causes
- Symptoms
- Solutions

The context type ENTRY has the parameters:

- type
- production status

The context type LOCATION has the parameters among others:

- Station ID
- Pressure
- Quantity
- Carbon monoxide concentration
- Methane concentration
- Velocity
- Resistance
- Roof
- Devices
- Events
- Obstructions
- Inby relation
- Outby relation
- Split relation

Taken in the aggregate, the context types and their associated properties completely describe the domain, and provide the necessary building blocks to reason about the domain. Before any additional discussion on the parameters, some properties of all parameters will be noted, specifically the "type," "expects," and "prompting-option" properties.

One of three values for the "type" property of parameters is allowed: "yes-no," "single-valued," and "multiple-valued." The "expects" property defines the allowable values of the parameter; "NUMBER" or a list of specific numbers, a list of specific entities, or "ANY" are examples of acceptable definitions for the "expects" property. The "prompting-option" property gives the system builder control over the means by which the expert system arrives at values for each parameter. For example this property can be set so that the system always infers the value, or always asks the user for the value, or asks the user for verification of an inferred value, or only asks the user after a failure to definitely establish a value through inference.

Table 1 lists examples of parameters and certain properties.

Table 1. Examples of Property Specification for Parameters.

<u>PARAMETER</u>	<u>Property</u>		
	<u>TYPE</u>	<u>EXPECTS</u>	<u>PROMPTING OPTION</u>
Monitoring system	single-valued	OK, Not OK, Questionable	Inferred (if possible)
Significant Problem	multiple-valued	Inadequate Quantity, Roof Fall, Excess Methane, Fire	Inferred (always)
Pressure Measure- ment	single-valued	Any Number	User ¹
			(Cont'd.)

Table 1. Examples of Property Specification for Parameters (Cont'd).

PARAMETER	Property		
	TYPE	EXPECTS	PROMPTING OPTION
Event	single-valued	Roof Fall, Equipment Move	Verify
Production Status	single-valued	Abandoned, idle, full	User
Note 1. In the future system, measured values will be obtained directly from the monitoring system.			

Another property of each parameter, known as "main-prop," allows the system builder to determine when the values for certain parameters will be obtained. This is an expedient feature since the system can be forced to first obtain values for fundamental parameters before extensive tree growth has occurred.

An examination of parameters, and their property specifications can provide insight into the workings of the domain. For the purposes of this paper such an examination is unwarranted, but some general comments about the structural component of the knowledge base may be useful. First, it is important to describe the domain with the least number of context types as possible, particularly in a small prototype system. As the number of context types increases, tree growth (dynamic context tree) increases geometrically, and user confusion increases exponentially. Memory size and increased execution time can also be problematic. Second, the EMYCIN inference structure is a tree, not a network. Accordingly, it is necessary that contextual relationships be "one to many" and not "many to one." This sometimes presented a problem in structuring the Ventilation Expert, but in all cases appropriate kludges were developed, with only a small decrease in utility. Third, the capability to control how the expert system obtains information is quite useful. It makes little sense to burden the system with complex inferences which the user could easily provide. Similarly, the system should not bother the user for information which could be easily inferred, especially in light of the fact that the user may provide erroneous information.

Problem Solving Component of the Knowledge Base

The rule base represents the knowledge, including the problem-solving expertise of the human expert. Development of the rule base requires an examination of three general aspects of the knowledge:

- types,
- representation, and
- acquisition.

Other specific issues are then associated with each of these aspects. The first two general aspects are examined in more detail.

The types of knowledge available to the Ventilation Expert are:

- Statutory requirements, e.g., required air quantities and velocities at specific locations, allowable practices, etc.;
- Ventilation plan, MSHA approved plan for operation of the ventilation system at the specific mine;
- Analytical models, e.g., fan laws, Atkinson's law, etc.
- Computer simulations, e.g., iterative routines such as PSU Mine Ventilation Simulator (MVS);
- Sensor measurements, i.e. inputs from the monitoring system;
- Heuristics, e.g. operational "rules-of-thumb," judgement, etc.

Except for the heuristical information, it is a straightforward, if somewhat tedious, task to represent these knowledge types. Also, obtaining knowledge from computer simulations represents a special problem which will be discussed later.

There are essentially three ways of representing knowledge in mini-MYCIN; they are rules of the form;

- causal,
- self-referencing, and
- meta-level.

Additionally certain rules can be streamlined using knowledge tables.

Causal rules take the form,

If	premise 1, and premise 2, and premise n
Then	conclusion 1, and conclusion 2, and conclusion.

The number of premises does not have to match the number of conclusions. If each premise is true, then the conclusions are affirmatively asserted; otherwise the conclusions are not asserted. The logical "else" operator may be used in premises or conclusions, as may be the logical "or" operator. An example of a causal rule from the Ventilation Expert is (parameters are in upper case):

If	SIGNIFICANT PROBLEM Inadequate Quantity is suspected	and
	$Q_{meas} < Q_{threshold}$	and
	FANS STATUS is definite	
Then	MONITORING SYSTEM NOT OK is suspected.	

Self-referencing rules are used to improve or reduce the certainty of a piece of knowledge in light of another piece of knowledge. For example it may be known a priori that a roof fall will obstruct an air

course, thereby reducing the quantity of air moving past the fall. If reduced flow is certain, but the cause is unknown, an expert would select a specific cause out of several, based on prior experience. In this case, roof fall as the cause, would be ranked as low probability. On the other hand, if the expert was told that the roof in the vicinity of the detected flow problem was an incompetent shale with closely spaced joints, then the probability of roof fall as the cause would be dramatically increased. Without getting into the probability aspects at this point, one example of a self-referencing rule is:

```
If      EVENT roof fall is suspected    and
      ROOF is jointed shale
Then    EVENT roof fall is definite.
```

Self-referencing rules are said to reason and conclude about themselves.

Meta-level rules are currently not used in the Ventilation Expert. This rule formulation allows strategic information to be imbedded in the rule base. For example given the presence of certain knowledge, the problem solving strategy could be modified in some fashion.

A backward chaining system such as mini-MYCIN requires a goal rule which represents the starting point of the expert system at execution time. The goal rule for the Ventilation Expert is:

```
If      there is a SIGNIFICANT PROBLEM, and
      the CAUSE is known
Then    find a SOLUTION.
```

The evaluation of this goal rule can be summarized without examining the reasoning network which will result from the evaluation.

First, the system will try to determine that a significant problem does exist. To do this it will look for all rules which conclude that a significant problem exists. Each of these rules will be evaluated until the premise, that a significant problem exists, is established, or until the relevant rules are exhausted. When attempting to evaluate a rule which concludes that a significant problem exists, the system will have to establish the truth of that rule's premise. In so doing, it will look for other rules which conclude about the truth of the premise in question. This backward chaining will continue until the validity of the original premise, in this case that a significant problem does exist, can be established. If this premise can be established as true, then the system will attempt to establish the truth of the next premise in the goal rule in a similar fashion. Otherwise execution is terminated. Once all premises have been satisfied the conclusions will be asserted. In the case of the goal rule, the conclusion "find a solution," will require additional reasoning.

One observation which can be made, based on the foregoing, is premises should be ordered with the least probable one given first down to the most probable one given last. Thus excessive time will not be wasted on rules which have premises which will seldom be true.

The foregoing discussion on rules has not considered the certainty with which premises or values are known, or to which conclusions can be asserted. The reasoning of human experts utilizes qualitative measures of certainty such as likely or more likely, and so forth. This presents somewhat a challenge to expert-system builders, since a more quantitative description of certainty provides a more familiar basis for analysis. However much theory has been developed to support reasoning under uncertainty.

EMYCIN utilizes a certainty factor concept which has been found to be satisfactory for the prototype Ventilation Expert. The certainty factor is associated with each instance of a parameter and reflects the degree to which the value assigned to that parameter is believed to be true (or false). The certainty factor can range from -1, corresponding to definite disbelief, to +1, corresponding to definite belief. Additionally the MYCIN developers formulated a conversion scale between the qualitative descriptors, e.g. not likely, and their numeric certain factor.

During the reasoning process the certainty factors which are propagated are not combined in some ad hoc fashion. Rather combination occurs through a carefully defined and tested set of procedures. Essentially measures of belief (MB) and disbelief (MD), for each hypothesis are updated each time a new piece of evidence is acquired. After all of the relevant rules have been evaluated a cumulative certainty factor is computed for each hypothesis by subtracting MD from MB. The hypothesis with the highest cumulative certainty factor is then asserted. Additional information on the concept and the procedures is found in reference 4.

The knowledge table is a useful tool to eliminate multiple articulation of certain knowledge. For example the following rule,

```
If  it is definite that the SENSOR ID is nnn
Then it is definite that the OUTBY RELATION is xxx, and
    it is definite that the INBY RELATION is yyy, and
    it is definite that the SPLIT RELATION is zzz,
```

would have to be rewritten for each possible value of SENSOR ID. However, by using a knowledge table, the foregoing rule could be written as,

```
IF  the value of SENSOR ID is known with certainty
```

```
Then it is definite that these parameters-OUTBY RELATION, INBY
RELATION, SPLIT RELATION should be transferred from the
knowledge table,
```

thereby allowing for a more efficient formulation of the rule base.

IMPLEMENTATION OF THE VENTILATION EXPERT

The knowledge-based expert system tool known as mini-MYCIN is written in NIL (New Implementation of Lisp) and will operate on a Vax in a VMS environment. The Ventilation Expert is currently being used on a microVAX 1, in a microVMS environment which has been tuned for NIL. The system is running with 2 megabytes of memory and 30 Mb of disk. This is approximately the smallest system configuration which will support mini-MYCIN. In fact, an additional 2 megabytes of memory is very desirable to reduce excessive paging and to increase the speed of the system. It would also be desirable to increase the disk space by at least another 20 Mb. However, as implied before, the important issue is building a small prototype system as early in the design process as possible, even with marginal facilities.

SUMMARY OF EXPERIENCES WITH THE VENTILATION EXPERT

Based on the experience derived from developing the Ventilation Expert certain observations can be made which may be useful to other domain experts contemplating the development of an expert system. One general observation is that it is very worthwhile to build a small prototype, even if compromises must be made. This effort will provide a clearer understanding of what is really needed, than would be possible otherwise. Despite compromises and shortcomings, mini-MYCIN was a very useful tool. Some other specific comments follow, primarily related to shortcomings which will (hopefully) be overcome in the next Ventilation Expert.

The backward chaining structure was useful for determining causes of problems and finding solutions, but quite cumbersome for determining if there was a significant problem. In the "real world" the monitored information would be checked for alarm states (values too high or low). If a sensor went to an alarm condition the Ventilation Expert would be invoked. If forward chaining were used, the expert system would know not only that a problem existed, but also where it existed. The reasoning network could then be developed using this initial knowledge. Currently, the control structure instantiates the location context for every sensor, which results in a massive tree with the attendant space and time problems. In a future system forward chaining should be used to establish the existence of a significant problem, and backward chaining should be used to determine the cause and find a solution.

The use of an expert system with a real-time monitoring system requires an expert system control structure which is cyclic. EMYCIN is not. It operates on one snapshot in time, and then must be restarted from scratch. This was not a serious problem for the prototype system, but would have to be resolved for an operational system. Ventilator Manager, which is a real-time version of MYCIN (5), demonstrates the feasibility of a real-time and cyclic system.

A more serious problem which plagued the development of the Ventilation Expert was EMYCIN's inability to work with spatial

relationships. The ventilation system of a coal mine is a complex network of aircourses. During the reasoning process it is often necessary to consider the variation of events at different points within the ventilation network. There was no straightforward way to do this, so a kludge using inby, outby, and split relationships was born, to permit inferences based on events at spatially related points. At this time an alternative does not seem to be available, although the use of a semantic network structure appears promising.

The nature of the reasoning process within the coal mine ventilation system domain frequently requires numerical calculations or the use of computer simulators. Numerical calculations might be used to compute trend information or determine the amount of proportional control to assert on a regulator, for instance. A simulator can be used to verify a hypothesis or predict the impact of a proposed solution for certain situations. Currently the calculations are handled with user defined functions (written in NIL); a software interface is being considered to allow information to pass from the Ventilation Expert to the FORtran-based simulator, MVS.

The use of the blackboard control structure with several independent knowledge sources is also being considered as a possible solution to the aforementioned problems. This seems particularly attractive in light of our conclusion that we seem to need a different type of expert system structure for each of the different aspects of the Ventilation Expert!

ACKNOWLEDGEMENT

The author gratefully acknowledges the patience, advice, and support provided by Prof. Peter Szolovits of the MIT Laboratory for Computer Science.

REFERENCES

1. USDI, Bureau of Mines, "Conceptual Investigation of a Management Information System for Underground Coal Mines," Contract No. J0348005, 10/31/83 - 1/30/86.
2. Ramani, R.V., Frantz, R.L., and R. Bhaskar, "An Analysis of the Decision-Making Processes in Mine Ventilation System Management," Internal Report prepared as part of BOM Contract J0348005 (Ref. 1 above), March 1985, 67 pgs.
3. Hayes-Roth, F.H., Waterman, D.A., and D.B. Lenat, ed., Building Expert Systems, Addison-Wesley Publishing Company, Inc., 1983, pp. 13-56.
4. Buchanan, B.G., and E.H. Shortliffe, ed., Rule-Based Expert Systems, Addison-Wesley Publishing Company, Inc., 1984, pp. 209-294.
5. Fagan, L.M., VM: Representing Time-Dependent Relations in a Clinical Setting, Ph.D. Thesis, Stanford University, 1980.

EMERGENCY MANAGEMENT DECISIONS AT THE LOCAL LEVEL

Robert S. Wilkerson
Federal Emergency Management Agency
Washington, D.C.

We are today often concerned with expert systems and systems in general. We are constantly concerned with the general question of the application of computers to the emergency management field.

I think it's very important we consider for a moment, who are the experts and what expertise we're proposing to integrate into such systems.

Generally, I feel we will find that emergency management experts fall into two categories: 1) those people who are experts because of their experience, generally in a limited number of instances. There are probably three or four clients of yours who have direct experience with an emergency situation and have developed some degree of expertise because of that; and, 2) those of us who are emergency management professionals at some other level who basically study that first group. We are the anthropologists of the emergency management world; the people who look at how others have made their decisions. I hope I'm already getting a message across: the proper focus of expert systems applications is at the local level.

Two years ago, I was asked to talk about the application of science and technology to State and local emergency management operations and decision-making. That venerable discussion in the halls of Congress deliberated the application of third generation computers, remote sensing devices, and telemetry devices, to a trade or profession which was operating at that point in the "technology of a legal pad." That situation has changed somewhat. I also argued that there was a problem particularly in emergency operations with the practitioners' receptivity to science and technology. There was difficulty in understanding how such tools might be applied. I have recently

become aware that this situation has changed dramatically. Let me give you an example of the change. I think it bears consideration during your deliberations as to when to apply expert systems, how long will it take, and what's the true market place.

Over the last several months we've asked a number of respondents about their application of computers, particularly to the evacuation problem. We asked a series of questions, hierarchical in nature, as to whether or not they had computers or had access to computers. Did they in fact use those computers in planning? Did they use those computers in operations? All the outputs are very interesting, but I want to quote one set of statistics that I found amazing. Of the respondents who said: a) that they didn't have a computer; b) that they had no current interest or intention of trying to access a computer for use in a planning and operations environment: almost half of them answered that they were interested in attending training or discussions on how the computer or technology might be applied to their situation. That represents to me a dramatic turnaround in the receptivity for the type of tools that you are discussing.

But what are we trying to apply them to? I said a "legal pad." What is the local environment? I think it is very important that we as Federal employees or you as technical experts always keep mind: there is no local environment. Now local government people, I didn't just sell the farm to the Feds. What I'm saying is that you can not homogenize the local government situation. You can't design an expert system in my opinion, one system that will apply to the legal structure and the decision making structure that exists in all of the States, and certainly can't design one that fits the situation in all of the local governments.

Let me deviate from the background of computer science into political science. We have a system in this country of strong governors and weak governors. We have at least three distinct groupings of government form at the local level.

The people who make decisions and therefore, would potentially use those expert systems are generally a group, not an individual, and they have varying powers and varying values. I think you have to bear that in mind.

But how do those decisions get made? Let's look at it first from an individual decision process. Individual expertise is often characterized as a person who can arrive at a decision, a sound decision, but can not generally explain precisely the variables considered.

That's especially true at the local government level where a wide variety of values are brought closer to the problems and decisions. At the local level, those decisions-makers, those users of the expert systems and also the experts, change frequently. Design of a static expert system, one that does not allow for frequent adjustment values for application to emergency management by the local government may not be cost effective.

We said we didn't have a homogeneous local government situation. You also can't consider emergencies as if they are homogeneous. The expert systems and the speed with which they need to be applied will vary based on the characteristics of the emergency. There's a lot more time to access expert systems in a hurricane than there is during a hazardous material emergency. One is a slow developing, fairly long forecast, even fairly predictable as opposed to the others which happens--now! Even hazardous materials incidents vary widely. A train derailment is dramatically different from a tank truck spill.

Experts systems can assist local decision-makers in making decisions ahead of time, and forecasting and outlining parameters and values. They can also provide interactivity and branching to other areas of expertise during emergencies. There is no reason why you shouldn't be able to ask a question then be branched elsewhere to expertise; or set your community decision parameters and recognizing a norm, from the decisions that a

collaborative group of experts have said they think they would make under similar situations, expect at least a flag saying to him; "you've gone outside those parameters; deviated from the norm; have you considered why you're doing this?"

The local government official's situation during an emergency, is a rags to riches information situation. There's either no information or there's more data than you can deal with. Clearly the expert system has to structure the information, has to assist the emergency manager in: 1) dealing with decisions when he has too much data, and 2) accessing sources of information when there is none.

We talked about the question of whether or not local governments are ready to use computers; whether they're ready for expert systems. I'd like to make a blanket statement: they are ready for the use of extensive technological tools that are presented to them in an understandable form, if they allow for the input of their values, if they allow for the sharing of data bases, if they do not create a command structure and remote decision making process. In short, they are ready for decision-making assistance tools. They are not ready for decision makers that are computers.

We sometimes talk about a five-year development process. The market for these types of systems is not five years down the road for two reasons: 1) the demand is now, and 2) there is a need to continue the process of building comfort with these complicated and sophisticated tools.

We can start today with basic decision trees that are constructed on an individual local government basis, forming the foundation for amalgamation into a larger decision or expertise sharing system. That market place is there today; the receptivity is there today; and we need to start. Cost benefit is going to be a continuing issue in this process. True cost

benefit will be found where the research and the development process renders a benefit today, incorporating user needs tomorrow, all as the product is being developed.

Receptivity to computers, involves incremental adjustment to the capabilities of a machine. A word processor replacing a typewriter has shown some resistance on the part of the user in corporation after corporation. Dropping an expert system without forewarning, without training, and without the experience that will come from participation in development, on an emergency manager will result in but one thing: an expert system that stays in the closet. It may "played with" on a day-to-day basis, but certainly in the heat of the moment, unless he, she, and they as a group are familiar with it, it will not be effectively used.

Simple decision assistance tools, the clear message to local government that what you're trying to provide is a tool, will not become a remote decision maker, cannot become a remote decision maker by the sheer question of legal authority. If we're going to talk about the bulk of emergency management situations, you are not talking about a national command structure. Thus, the system has to be distributed, decision parameters have to be put in at the local level, and the data has to be shared. There is also a need within the national command structure. Both local and national needs can be supported. The development of the data base and a growing familiarity with the use of expert systems will build a tool that will fit the national emergency environment and yet at the same time provide a cost effective capability for day-to-day emergency management use.

Three or four strong points; we're talking about a group user not an individual; we're talking about masses of group users not a single group. We can develop a decision support system; an expert system to use at the national level; but without the local input, without the local parameters,

it will be of limited, if any, use. So what I'm advocating is a building block concept that starts at the local level and builds up to the national level. If you will, an expert system that's based on a multitude of experts as opposed to a select few. We've got to incorporate the experience and the risk parameters of a number of people.

Are we talking about a centralized system? If the definition for a centralized system is a main frame structure, into which every local user must go each time then the answer is no. If we're talking about a centralized system, better defined as an interconnected system which ties a number of experts together, a number of users together, then we are talking about a centralized system. We desperately need to move from the individual experiential mode into that experience sharing process, possible through interconnectivity.

If the system is structured so that it interacts with the decision-maker in a manner that asks questions that lead to deliberate answers and then analyzes the answer against what would have been the expert system recommendation, it will be better received than if it simply provides an answer and even gives a statistical value of the reliability of that answer. Ask the question, "What would you do? You said you would consider these parameters, bam, bam, bam. Have you considered X, Y, and Z? What is your decision? Your decision is somewhat inconsistent with what your preprogrammed answer was. It is also inconsistent with the answer we provided."

"Is that your final decision?" Yes? Then the system learns from that and considers it subsequently in the life of an emergency when the same type of decision is called for again. In that situation, you deviated from your standard answer, you deviated from the past solution, and you deviated from the expert system's answer.

"Do you still want to do this?" Yes! Situations are going to change suddenly. Decision factors that are built into the system; because of fatigue, stress, or public pressure will cause the expert, the decision-maker, to learn during the process. Therefore, you have to continue to ask questions because he's under stress and you have to make sure that he or she doesn't forget. The system becomes an automated tickler system that then compares and reminds the decision-maker that he has varied somewhat, reminds him to consider the justification. Why? One of the biggest single issues in public safety today is justification after the fact. Why did you make that decision? Where are your records? What is your liability? Did you follow your plan, San Jose, California? No. Why? Because the information I had (which I plugged into the expert system which gave me a record of the decision parameters and subsequently the decision) varied somewhat from what I expected in my plan. The decision tool, your expert system and the software and hardware that makes up that system, serves a multiple purpose to the local government users.

Remember who the user is. Expert systems are critically needed in emergency management and they're needed now. But if you forget who the ultimate user is, we will repeat past failures. A project in Florida involving NOAA, FEMA, and the State of Florida and over 30 local governments, the "Probabilistic Decision Tool for Hurricane Evacuation", left out several critical success parameters.

- 1) the expert system was essentially based on the decisions of one expert;
- 2) the only interactivity of the expert system cross-referenced to that one expert.
- 3) no local parameters for the risk was acceptable, no local definition of what actions they thought should take place under certain conditions.

Make these mistakes nationally, and it will take five to ten years to develop a meaningful expert system.

One final comment on the interconnectivity standpoint. When I preach distributive systems, I don't want to come across as saying that it is not critical that one local government be able to learn from another local government. So if we are to go with a micro- or PC-based expert system, we must rapidly move to a situation that a flag is placed in the decision process of City A, based on the recent experiences of City B. We can no longer afford to spend years transferring experience from one local government to another.

The SMOKEY Project: An Intelligent Sensor-Based Fireman's Assistant

Jaime G. Carbonell

Carnegie-Mellon University

Pittsburgh PA, 15213

28 April 1986

Abstract

This paper reports on the SMOKEY project, an artificial intelligence (AI) approach to automating the diagnosis, assessment and suppression of fires in large man-made structures such as aircraft carriers or building complexes. The requirements for automation and its desirability in certain classes of emergency situations are discussed. Based on these requirements, the design and implementation of SMOKEY, a successful AI prototype, emerged. Certain computational mechanisms were employed and extended, including agenda-based control structures for dynamic task prioritization in time-critical situations, and evidentiary combination rules for reasoning under uncertainty.

The task-domain research was supported by the Naval Air Systems Command under contract number N60530-83-M-487A, and the research on the underlying AI technology was supported in part by the Office of Naval Research under contract number N00014-84-K-0345.

1. Background: The Need for Automated Fire Assessment

The detection, assessment and management of accidental fires on man-made structures, such as ships and buildings, requires rapid data-gathering and decision making. Present sensor-based equipment, such as smoke detectors connected to fire alarms, provide adequate technology for the detection phase. But, assessing the severity and nature of the fire requires the presence of skilled personnel. Moreover, formulating an effective containment or suppression strategy requires an accurate and timely assessment, prediction of expected spread or growth of the fire, knowledge of available fire-fighting equipment, and extrapolation of expected sideeffects of the fire (e.g., structural weakening or release of toxic fumes). Typically, skilled human intervention is required to make real-time decisions beyond the initial detection stage.

There are many situations, however, where automating the assessment of the severity and possible spread of the fire, as well as automating the selection of an appropriate fire-fighting strategy could prove far more effective than present practice requiring human intervention. Rudimentary systems, such as sprinklers and the release of heavy inert oxygen-suppressant gases, tied directly to smoke or heat sensors are in existence and undergoing continued improvement. However, here we focus on much more sophisticated systems providing flexible functionality comparable to that of skilled fire-fighting personnel.

Automating the fire assessment and strategy selection is most appropriate in situations where too much information overwhelms human judgement, either because human reaction time is too slow, or because the requisite expertise is unavailable or too far removed from the site. We consider each case in turn:

- In some situations accurate decision making requires processing vast amounts of data at a rate significantly higher than human capacity permits. Such a situation exists, for instance, aboard aircraft carriers where the information from thousands of sensors converges upon a complex panel, and the officer in charge must make decisions in seconds based not only upon the dynamically changing sensor readings, but also upon topological knowledge of the ship (which rooms are connected to which others; what

combustible materials are stored where; which passages can be chosen as escape routes for anyone trapped by the smoke; what fire fighting equipment is stored where; etc.) The complexity of an aircraft carrier coupled with the real-time nature of the task makes it difficult for a human to make accurate decisions in a timely manner, especially for large fires or multiple simultaneous emergencies.

- There are certain situations, where it is not the volume of relevant data that overwhelms human capacity, but the need for split-second reaction time prevents human consideration, let alone accurate judgement. For instance, a fire rapidly spreading towards a fuel tank or towards the magazine, requires immediate action (e.g., flooding the fire with water or inert gas, or closing fuel line valves). The cost of inaction or delayed action might be far greater than the benefits derived from more considered deliberation. On the other hand, flooding the library or the computer room with water, is likely to cause more damage than the fire. Therefore, intelligent assessment is required, and if the danger is immediate action should be taken by the automated system. If not, available human expertise should be consulted, or at least provided the option of overruling the recommendations of an automated system.
- There are situations where assessment of a fire situation requires on-site expertise, but where that expertise simply cannot be present in human form just for the unlikely eventuality of a fire. Most large commercial buildings, hotels, apartment houses, factories, etc. are wired with sophisticated smoke and heat sensors, but cannot have an everpresent on-site fireman in the eventuality that a fire may potentially occur. In such cases, automating the expertise in a central system connected to the sensors can result in timely assessments and calls to the fire stations with information as to the location, spread direction, combustion materials (and therefore required fire-fighting equipment), occupancy and smoke-clear escape routes. Such information can result in more timely and appropriate actions by the human firefighters (or by the automated system lighting up the clear escape routes, informing the arriving firefighting personnel of the state of the combustion, and perhaps controlling the appropriate sprinkler or inert gas system).

2. SMOKEY: The Task Definition

The task domain for the SMOKEY project is fire assessment and fire control aboard modern aircraft carriers such as the USS Carl Vinson, and thus the research focused primarily on the first of the three scenarios above: real-time processing of large volumes of time-variant information. Although we believe the basic technology is extensible to the other two situations as well, the rest of this paper focuses only on the SMOKEY analysis and its prototype implementation.

In order to test the feasibility of automating the sensor based fire assessment tasks for an aircraft carrier, we started by interviewing several naval officers with operational experience in fire assessment and suppression processes. We recreated some scenarios where different classes of fires might be manifest, and recorded their step-by-step reasoning processes. In order to get more accurate assessments and more accurate reporting, we slowed down the simulated development of the fire by a factor of 10 in several cases. Although the details of the knowledge engineering process are beyond the scope of this paper, we mention that our technique combined classical protocol analysis, with post-hoc interrogation to re-analyze and justify past decisions. Afterwards, we reduced these problem-solving traces to rules and augmented those rules with algorithmic processes for graph traversal tasks, as discussed below, and with evidentiary combination rules to replace the ad-hoc methods we observed for reasoning under uncertainty. Finally, we hypothesized the existence of additional sensors (mostly differential air pressure sensors) that could easily be installed. We then constructed a computational prototype to investigate the feasibility of a fully automated fire diagnosis, assessment and advice generator, based on simulated sensor readings.

3. Design Requirements

SMOKEY was specifically designed to address problems of real-time multi-sensor fusion using AI technology. There are few AI systems that perform real-time analysis of external sensor reading, one such being SUX [6] for analysis of acoustic signatures from submarines. However most AI systems, even interactive ones such as XSEL [5] or POLITICS [3, 1] do not address serious real-time demands. As such, the development of SMOKEY dictated an analysis of its requirements and a significant effort

in the extension and integration of AI-based computational techniques. The first step was to build a prototype to demonstrate the feasibility of real-time automation of decision making as a fire develops.

Our design criteria were to develop a system to perform the following tasks:

- Take readings from large numbers of sensors -- putting some 3,000 sensors of different types including smoke detectors, heat detectors, air-flow meters (to determine air circulation) position sensors (to determine whether hatches or doors are open or closed), and more complex sensors, such as smoke analyzers (to determine whether the smoke contains ozone, hydrocarbons, metallic oxides, etc.) -- with the assumption that the bulk of the sensors are simple smoke detectors, and the other, more sophisticated sensors could be assembled and installed from commercially available components.
- Localize and assess the extent and nature of the fire. The old adage "where there's smoke there's probably fire" is too simplistic. The air circulation system aboard ship, or in a modern building, quickly distributes the smoke well beyond the immediate combustion area.
- Assess the criticality of the fire. Criticality depends on many factors, including:
 - The nature of the combustion by-products. Certain plastics, for instance, produce highly toxic fumes, whereas paper fires do not.
 - The nature of the combustion materials. Oil fires are harder to put out than paper fires. And, metallic combustion (such as magnesium) is harder yet to control.
 - The location and direction of spread of the fire. A fire spreading in the direction of the magazine or the jet fuel tanks is obviously more serious than one of similar size in a remote galley or lavatory.
 - The present extent of a fire, including differentiation between previously affected areas and the combustion front.
 - The amount of combustible material in the immediate area that can feed the fire.
 - The accessibility of the affected region to firemen and fire suppressant materials.

- The extent to which the fire blocks normal exit routes or otherwise hampers escape.
- The degree to which external factors impact fire suppression (low water pressure, untrained crews, etc.)
- Take direct on-line action), such as lighting up signs to signal exit passages for endangered personnel that avoid the smoke afflicted areas (as opposed to lighting up the shortest escape routes which might lead to impassible regions), sounding appropriate alarms, notifying appropriate personnel, etc....
- Recommend off line actions to the fire fighting crew, including what fire suppressant equipment they should use (e.g. water or oxygen suppressants such as foam or CO₂), and what additional equipment is necessary (e.g. gas masks if toxic fumes are present).
- Perform all of the above in real time with partial information -- and be able to reconsider past decisions in light of new information. In most real-time decision-making tasks, conclusions must be reached using available data, rather than waiting until all relevant information is gathered and analyzed. A perfect analysis of what should have been done to combat a fire hours earlier is of much less use than timely if not always perfect recommendations.
- Allow for errorful sensor inputs caused by random sensor failures, and possibly by the fire itself burning out sensors or melting the wires leading from sensors inside the combustion front.

Given the magnitude of the task in trying to achieve all of the objectives set forth above, we opted for a comprehensive design and the implementation of a small pilot system to investigate the feasibility of the endeavor. Our first step was to analyze computational mechanisms required to cope with all the complexities in real time. SMOKEY was the resultant successful prototype converged upon, after two earlier less successful attempts were abandoned in the design stage. SMOKEY runs on a dedicated LISP workstation (presently the microvax II), but to place it in operational use would require its translation to a conventional language like C (for faster operation and for direct interfacing

with all the external sensors).

4. The SMOKEY Architecture

SMOKEY requires a static knowledge base and a dynamic working memory. The former is quite large and consists of:

- A graph structure encoding the topology of the 3D structure (a hypothetical part of an aircraft carrier in our test domain -- but could be replaced by a building complex or any other structure containing rooms, passages, doors, and air circulation systems). The topological graph represents rooms as structured nodes (with which it associates the placement of sensors, and the potentially flammable materials contained therein), and represents passages and air vent connections as labeled arcs connecting the nodes. This graph was implemented in FrameKit [2], a simple and fast frame-based language for storing static information economically by inheriting general facts shared by multiple instances.
- Information about flammable substances and their combustion characteristics, including gaseous byproducts, flash points, etc.
- Information about fire suppressant substances and their means of delivery. Side-effects of these substances are also recorded, such as the damage water can produce to books or computer equipment.

The dynamic knowledge base consists primarily of a time-stamped sensor-poll buffer indicating the last 3 readings of each sensor and the time at which it was read (thus trends can be detected -- it may be important, for instance, to determine if a room is getting hotter or cooler after the sprinkler system goes off to infer the effectiveness of the fire suppression method). The second component of the dynamic memory is a set of pending hypotheses -- the system's current beliefs about the nature of the fire prior to recommending action -- rank ordered by their criticality (thus decisions about whether the fire is about to spread towards the magazine are given immediate attention, whereas hypotheses about water damage to documents is given less priority).

5. The Agenda Control Structure

In order to control the multiple processes (sensing, generating hypotheses, testing hypotheses, suggesting actions, updating working memory from new information gathered, etc.) and provide real-time response, a computational discipline is required in which one can assign priorities to each task dynamically depending upon what problems are judged most critical at that time. These priorities determine the focus of attention of the SMOKEY interpreter -- with low priority tasks being attended if and only if there are no higher priority ones awaiting action. We selected an extended agenda-control mechanism implemented in RuleKit [9], a rule-based language we developed precisely for real-time control tasks.

The agenda mechanism consists of an ordered set of priority levels (called "buckets"), each of which groups a list of tasks of roughly equivalent priority. A task can be:

- A conditional sensing action, e.g., polling external sensors and thus updating the sensor-poll buffer.
- A status update, posting and propagating values of sensor readings on the topological graph.
- A hypothesis formation, e.g., if a certain configuration of readings is detected, a hypothesis about the nature or spread of the fire is posted in the "pending hypotheses" list
- A hypothesis test, as more information is gathered past hypotheses are confirmed, ruled out, or modified.
- An immediate action generator, e.g., lighting up signs once a safe smoke-free exit passage is calculated.
- A recommendation generator, e.g. advising the fire-control team as to the best fire-suppressant substance to use given the assessed nature and extent of the fire.
- An internal attention focus shift, where reprioritization of tasks occurs as new information

is discovered (e.g. a new combustion front breaks out), or a past hypothesis that was guiding action is later invalidated, again in light of new information.

Priority-1:	Task _{1,1} ,	Task _{1,2} ,	...	Task _{1,n}
Priority-2:	Task _{2,1} ,	Task _{2,2} ,	...	Task _{1,m}
.
.
.
Priority-k:	Task _{k,1} ,	Task _{k,2} ,	...	Task _{k,r}

The agenda mechanism operates by attempting to perform each task at the first level of priority. If any task is carried out, all tasks at that level are again considered. However, if none are carried out (either because all are awaiting more information, or are awaiting external action such as the firemen arriving on the scene), then the tasks at the next priority level are considered, and if possible carried out. If any task at any level was carried out, control returns to the first priority level (new information gathered, or new hypotheses posted, or simply time having passed, may now enable a higher priority task to execute). However, if at any priority level no task is executable, control passes to the next lower priority level. In this fashion, high-priority tasks are always tried and if possible carried out, and low-priority ones are carried out only when awaiting for the high priority ones to be executable.

Each task in the agenda consists of five components:

TASK-i	TEST:	(check sensors, hypotheses, state, time, ...)
	ACTION:	(update hypothesis, state OR external action)
	ADD:	(insert new tasks in agenda at new priority)
	DEL:	(remove or demote other tasks in agenda)
	KTEST:	(check whether to remove this task from agenda)

The *test* field is a conjunction of predicates on the sensor poll buffer, the pending hypotheses, the clock, and any other internal working memory elements. If all these predicates are satisfied, the task is executed, which means that the set of internal and actions specified in the *action field* are carried out sequentially. Also, as a task is executed, new tasks may be inserted into the agenda at specified priority levels (specified in the *add* field), or old tasks may be removed (specified in the *del* field). For instance, if a task confirms the hypothesis that an electrical fire is present, its action can select the fire suppressant (e.g., CO₂), the add field inserts a task to stop the electrical current through the afflicted

wires, and the *del* field removes the tasks looking for combustion byproducts of other types of fires.

If the test field of a task is not satisfied, the *ktest* field is evaluated to determine if the task should remove itself from the agenda. Thus, tasks can have their own expiration conditions. Real time processing requires that only the potentially relevant tasks be considered at any one time -- a small fragment of the total set of tasks that may apply in different circumstances. The dynamic addition and deletion of tasks in the agenda, together with the prioritized task consideration provides the necessary software architecture to support real-time sensor analysis and assessment of time-variant phenomena.

6. Evidentiary Combination Methods

Any complex network of sensors cannot be expected to function perfectly all the time. Sensor malfunction or wiring problems will yield occasional false positives (spurious readings) and persistent false negatives (non-functional sensors). The artificial intelligence program performing diagnosis and situation assessment from these sensors cannot ignore such problems, and must combine with statistical techniques to provide robust behavior in the presence of uncertainty. We borrowed the *evidentiary combination* rules below from the POLITICS system [1] used to simulate human decision making in the face of possible uncertainty of data and possible uncertainty of the inference rules themselves. Other approaches such as those of MYCIN [10] or PROSPECTOR [4] were less well suited to the current task.

In essence, the evidentiary combination rules are derived from standard probability theory, assuming independence of possible observation errors (e.g. flawed sensors), and independence among the reliability measures on the inference rules. Such independence assumptions provide the most accurate simplification of otherwise intractably complex mathematics.

Sufficiency conditions are considered separately from necessity conditions for arriving at a given conclusion given a set of observations. Thus, the problem of calculating the certainty of a conclusion factors into a disjunctive component (only one sufficiency condition need be present with high

reliability), and a conjunctive one (all necessity conditions need be present with high reliability).

Looking at the disjunctive case first, we identify two sources of uncertainty, the *a-priori* uncertainty associated with the reliability of each observation (the $P(D_i)$'s in the formula below), and the degree to which an observable is a sufficient condition to assert by itself the conclusion (the $I(D_i, R_j)$'s in the formula below). For instance, detecting partially oxidized hydrocarbons is a weak indicator of oil fires (internal combustion motors may produce high ambient levels of such byproducts, and some plastic fires may also produce them), whereas detecting high concentrations of ozone is a much stronger indicator of electrical-arc fires. The I 's in the formula measure the strength of each indicator. Taking together the strength of an indicator with the reliability of the observation produces a measure of confidence in the conclusion. However, multiple measures must be combined, as multiple sources of corroboratory information can produce a much higher confidence in the conclusion, even if each source is individually somewhat suspect. The disjunctive combination formula below integrates multiple observation certainties and indicator strengths to yield a composite confidence measure on the conclusion.

$$P_{or}(C_j) = 1 - \prod_{i=1}^k (1 - P(D_i)I(D_i, R_j))$$

$$P_{\&}(C_j) = \prod_{i=1}^k [1 - I(D_i, R_j)(1 - P(D_i))]$$

The corresponding conjunctive formula for necessity conditions is given in the second equation. In both cases we assume that all $P(D_i)$'s are true probabilities, and that indicator strengths are normalized in the $[0,1]$ interval, with 0 meaning no correlation and 1 meaning absolute correlation. The output $P(C_j)$'s are then true probabilities (whose accuracy depends on the validity of the independence assumptions mentioned earlier). Inference rules may thus be chained and their certainty values propagated by applying the formulae above. We found these rules to be quite workable in SMOKEY, and much faster to compute than Dempster-Shafer [8], fuzzy logic [11] or other complex methods less suitable for real time performance. (The reader interested in probabilistic logic in AI is referred to [7] as a starting point.)

7. Concluding Remark

The SMOKEY project produced a successful pilot implementation of a fully-automated fire assessment and suppression system, but thus far has remained within the laboratory. It may now be time to build a scaled-up version and place it in a real operational setting as a field trial with complete manual overrides by human operators. The primary contributions of SMOKEY have the computational mechanisms developed for real-time artificial intelligence systems, which are already finding applications outside their initial context. RuleKit and FrameKit, for instance, have seen multiple uses at many different sites, and the agenda-driven discipline is gaining popularity. Finally, the computationally tractable evidentiary combination methods proved to be just what the fireman ordered for efficient real-time probabilistic reasoning. With all these tools in place, we are now ready to start engineering larger scale real-time decision making systems.

8. References

1. Carbonell, J. G., "The POLITICS Project: Subjective Reasoning in a Multi-Actor Planning Domain," in *Computer Science Research Review*, Carnegie-Mellon University, 1979-1980.
2. Carbonell, J. G. and Joseph, R., "The FrameKit⁺ Reference Manual", CMU Computer Science Department internal paper.
3. Carbonell, J. G., *Subjective Understanding: Computer Models of Belief Systems*, Ann Arbor, MI: UMI research press, 1981.
4. Duda, R. O., Hart, P. E., Konolige, K. and Reboh, R., "A Computer-Based Consultant for Mineral Exploration," Tech. report 6415, SRI, 1979.
5. McDermott, J., "XSEL: A Computer Salesperson's Assistant," in *Machine Intelligence 10*, Hayes, J., Michie, D. and Pao, Y-H., eds., Chichester UK: Ellis Horwood Ltd., 1982", pp. 325-337.
6. Nii, P. H. and Figenbaum, E. A., "Rule-Based Understanding of Signals," in *Pattern-Directed Inference Systems*, Waterman and Hayes-Roth, 1978, eds., New York: Academic Press, 1978.
7. Nilsson, N. J., "Probabilistic Logic," *Artificial Intelligence*, Vol. 28, No. 1, February 1986, pp. 71-87.
8. Shafer, G. A., *A Mathematical Theory of Evidence*, Princeton University Press, Princeton, NJ, 1979.
9. Shell, P. and Carbonell, J. G., "The RuleKit Reference Manual", CMU Computer Science Department internal paper.
10. Shortliffe, E., *Computer Based Medical Consultations: MYCIN*, New York: Elsevier, 1976.
11. Zadeh, L. A., "Fuzzy Logic and Approximate Thinking," *Synthese*, Vol. 30, 1975, pp. 407-428.

FIRST SYMPOSIUM ON THE THEORY AND APPLICATION
OF EXPERT SYSTEMS IN EMERGENCY MANAGEMENT OPERATIONS

- - - - -
Sponsored by

The Federal Emergency Management Agency
and
The National Bureau of Standards

- - - - -
Wednesday, April 24 and Thursday, April 25, 1985
- - - - -

Department of Commerce Main Auditorium
14th Street and Pennsylvania Ave., N.W.
Washington, D.C.

PROGRAM

Wednesday, April 24, 1985

9:00 A.M. - 9:10 A.M.

Opening Remarks..... Saul I. Gass
Symposium Chairman
University of Maryland
National Bureau of Standards

9:10 A.M. - 9:20 A.M.

Welcome..... Gen. Louis O. Giuffrida
Director
Federal Emergency Management
Agency

9:20 A.M. - 9:30 A.M.

Introductory Remarks..... Raymond G. Kammer
Deputy Director
National Bureau of Standards

9:30 A.M. - 10:30 A.M.

A Theory of Expert Systems..... A. B. Winston, Purdue
University

10:30 A.M. - 11:30 A.M.

The Management of Uncertainty in Expert Systems..... L. A. Zadeh, University of
California, Berkeley

LUNCH 11:30 A.M. - 1:00 P.M.

1:00 P.M. - 2:00 P.M.

New AI Tools for Building Management Aid Systems..... D. L. Waltz, Thinking
Machines Corporation

2:00 P.M. - 3:00 P.M.

TIMM: A Tool for Building Expert Systems..... J. N. Craig, General
Research Corporation

3:00 P.M. - 4:00 P.M.

Emergency Management Decisions at the Local Level.... R. S. Wilkerson, Federal
Emergency Management
Agency

4:00 P.M. - 5:00 P.M.

Decision Support Systems in Emergency Management
Operations: A Microcomputer Perspective..... F. De Balogh, University
of Southern California

Thursday, April 25, 1985

9:00 A.M. - 10:00 A.M.

Information Technology for Emergency Management..... R. L. Chartrand, Library
of Congress

10:00 A.M. - 11:00 A.M.

Emergency Management for Chemical Spills..... C. Johnson, Oak Ridge
National Laboratory

11:00 A.M. - 12:00 Noon

Roles of Simulation in the Application of Expert
Systems in Emergency Management Operations..... A. B. Clymer,
Clymer Technology

LUNCH 12:00 Noon - 1:30 P.M.

1:30 P.M. - 2:15 P.M.

Expert System for Fire Emergency Management..... G. N. Berlin,
Modeling Systems
Incorporated

2:15 P.M. - 3:00 P.M.

Sensor-Based Diagnosis in the SMOKEY System..... J. Carbonell,
Carnegie-Mellon
University

3:00 P.M. - 3:45 P.M.

The VENTILATION EXPERT for Underground Coal Mines..... J. L. Kohler,
The Pennsylvania
State University

3:45 P.M. - 5:00 P.M.

Panel Discussion

FIRST SYMPOSIUM ON THE THEORY AND APPLICATION OF
EXPERT SYSTEMS IN EMERGENCY MANAGEMENT OPERATIONS

Speakers' Names and Addresses

Dr. Carroll K. Johnson
Oak Ridge National Laboratory
Building 4500 North
Mail Stop C18
Oak Ridge, TN 37830

(615) 574-4975

Dr. David L. Waltz
Thinking Machines Corp.
245 First Street
Cambridge, MA 02142

(617) 876-1111

Mr. Geoffrey N. Berlin
Modeling Systems, Inc.
1718 Peach Tree St. N.W.
Atlanta, GA 30309

(404) 876-9977

Dr. Andrew Whinston
Krannet School of Management
Purdue University
West Lafayette, IN 47907

(317) 494-4468

Dr. Jaime Carbonell
Computer Science Department
Carnegie-Mellon University
Pittsburgh, PA 15213

(412) 268-3064

Mr. Robert Lee Chartrand
Congressional Research Service
Library of Congress
Washington, D.C. 20540

(202) 287-7056

Mr. Ben Clymer
Clymer Technology
32 Willow Drive, Apt. 1B
Ocean, NJ 07712

(201) 493-4364

Dr. Lotfi A. Zadeh
Department of Electrical Engineering
University of California
Berkeley, CA 94704

(415) 642-4959

Dr. Frank De Balogh
Decision Support Laboratory
Institute of Safety and Systems Management
Room 109
University of Southern California
Los Angeles, CA 90089-0021

(213) 743-4048

Dr. Jeffery L. Kohler
Department of Mineral Engineering
The Pennsylvania State University
University Park, PA 16802

(814) 865-4491

Mr. Joseph N. Craig
General Research Corp.
7655 Old Springhouse Rd.
McLean, VA 22102

(703) 893-5900

Mr. Robert S. Wilkerson
State and Local Programs and Support
Federal Emergency Management Agency
500 C Street, S.W.
Washington, D.C. 20472

(202) 646-2861

APPENDIX I: EXPERT SYSTEMS AND EMERGENCY MANAGEMENT: BACKGROUND PAPER

Saul I. Gass
College of Business and Management
University of Maryland
College Park, MD 20742

and

Robert E. Chapman
Mathematical Analysis Division
Center for Applied Mathematics
National Bureau of Standards
Gaithersburg, MD 20899

This background paper was prepared for the speakers at the Symposium. It served to introduce the disparate communities of expert systems and emergency management to the basic problems and procedures of the two fields. We reproduce it here as background material for those readers who are external to either or both of these fields.

EXPERT SYSTEMS AND EMERGENCY MANAGEMENT: BACKGROUND PAPER

The National Bureau of Standards (NBS) is conducting a research effort under the sponsorship of the Federal Emergency Management Agency (FEMA) to identify specific areas of emergency management operations where the use of expert systems would be of greatest benefit to FEMA. A symposium, held on April 24-25, 1985, is a key component of the NBS research effort, as it provided a forum through which researchers in the area of artificial intelligence can meet and interact with emergency management personnel. The symposium was of immediate assistance to FEMA because it provided guidance on how FEMA should structure and budget its expert systems development and training programs. The following material was prepared as background information for the Symposium's speakers.

I. An Introduction to Expert Systems

Although there are numerous texts and articles available on the subject of expert systems, one which is particularly useful as a reference document for the upcoming symposium is Building Expert Systems.¹ This text is recommended because chapter 10 contains a discussion of the Oak Ridge National Laboratory (ORNL) hazardous chemical spill program, an expert system which addresses an emergency management problem. The following brief description of an ideal expert system has been excerpted from the text; the system is illustrated schematically in Diagram 1.

At the most basic level, an expert system may be defined as a computer-based model which is able to replicate the decision of one or more human experts in a well-defined subject area. Furthermore, the types of problems addressed by expert systems are sufficiently difficult to require significant human expertise for their solution. An expert system is therefore a "model" of the expertise which the human "expert" brings to bear on the problem. The major components which serve to model a human decision maker's expertise are referred to as the knowledge base and the inference procedure.

The knowledge base records rules, facts, and information about the current problem that may be useful in formulating a solution. The inference procedure provides a control structure for bringing to bear the information contained in the knowledge base (i.e., facts and rules) on the problem.

Notice that it is the distinct nature (i.e., separation) of knowledge base and the inference procedure which taken together captures the human decision maker's expertise. This separation and its associated flexibility is absent in traditional models which use a rigidly defined control structure to manipulate the information contained in a database. Knowledge is therefore the key and to a great degree the performance of an expert system is a function of the size and quality of the knowledge base that it possesses. Some authors thus choose to use the term knowledge-based expert systems to emphasize the importance of knowledge in modeling human expertise.

¹F. Hayes-Roth, D. A. Waterman and D. B. Lenat (editors), Building Expert Systems, Reading, MA: Addison-Wesley, 1983.

An ideal expert system contains: (1) a language processor for problem-oriented communications between the user and the expert system; (2) a global data base for recording intermediate results; (3) a knowledge base comprising facts as well as planning and problem-solving rules; (4) a control structure, which includes an interpreter that applies the rules, a scheduler to control the order of rule processing, a consistency enforcer that adjusts previous conclusions when new data (or knowledge) alter their bases of support; and (5) a justifier that rationalizes and explains the system's behavior.

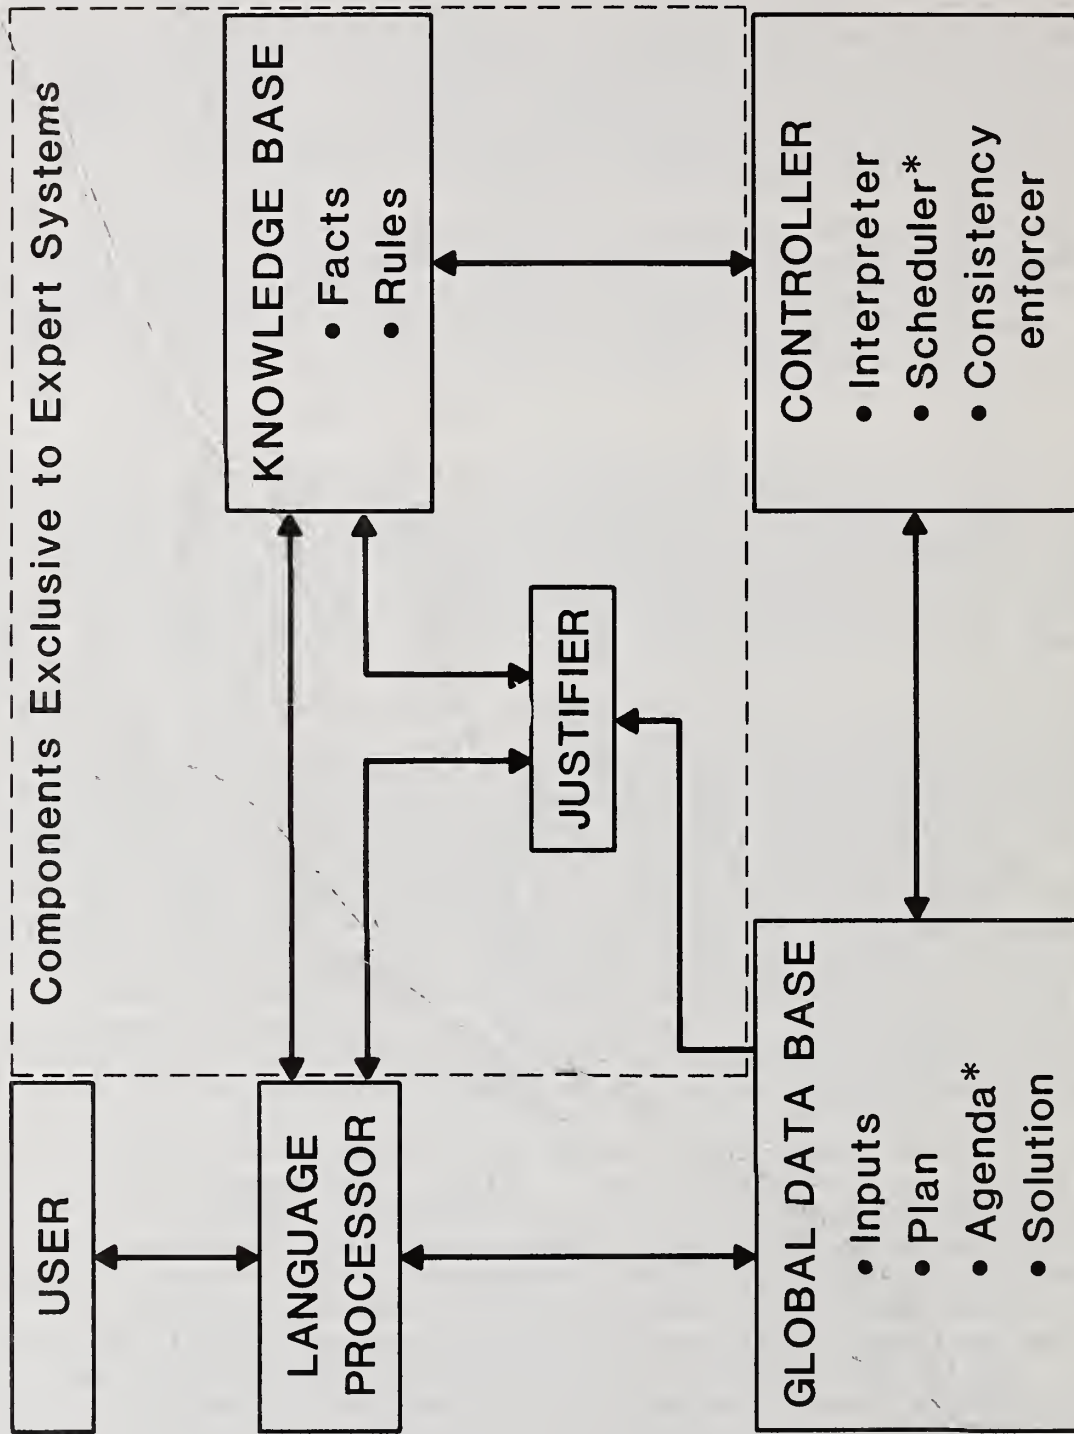
The user interacts with an expert system via a problem-oriented language. The language processor mediates information exchanges between the expert system and the user. Typically, the language processor interprets questions posed by the user, commands given by the user, and volunteered information. Conversely the language processor formats information generated by the system, including answers to questions, explanations and justifications for its behavior, and requests for data.

The knowledge base of an expert system consists of facts and rules. The "facts" constitute a body of information that is widely shared, publically available, and generally agreed upon by experts in the field. The "rules" are often based upon less well-known heuristics (rules of plausible reasoning, rules of good guessing) that characterize expert-level decision making in the subject area.

The global data base records the user's inputs, intermediate hypotheses, and decisions that the expert system manipulates. Every expert system uses some form of intermediate decision representation. However, only a few explicitly employ all of the components shown in Diagram 1. In addition to the user's inputs, three types of elements are identified as being included in the global data base: (1) plan; (2) agenda; and (3) solution elements. Plan elements describe the overall or general attack the system will pursue against the current problem, including current plans, goals, problem states, and contexts. The agenda elements record the potential actions awaiting execution, which generally correspond to knowledge base rules that seem relevant to some decision place in the global data base previously. The solution elements represent the candidate hypotheses and decisions the system has generated thus far, along with the dependencies that relate decisions to one another.

The control structure determines how to use the rules contained in the knowledge base; it contains three elements: (1) an interpreter; (2) a scheduler; and (3) a consistency enforcer. The interpreter executes the chosen agenda item by applying the corresponding knowledge base rules. Generally, the interpreter ensures the relevance conditions of the rule, binds variables in these conditions to particular global data base solution elements, and then makes those changes to the global data base that the rule prescribes. The scheduler maintains control of the agenda and determines which pending action should be executed next. Schedulers may embody considerable knowledge, such as doing the most profitable thing next. To apply such knowledge, the scheduler needs to give each agenda item a priority according to its relationship to the plan and other extant solution elements. To do this, the scheduler generally needs to estimate the effects of applying

Diagram 1. Anatomy of an Ideal Expert System



*These items are unique to expert systems

the potential rules. The consistency enforcer attempts to maintain a consistent representation of the merging solution. Most expert systems use some kind of numerical adjustment scheme to determine the degree of belief in each potential decision. Such schemes attempt to ensure that plausible conclusions are reached and inconsistent ones are avoided.

Finally, the justifier explains the actions of the system to the user. In general, it answers questions about why some conclusion was reached or why some alternative was rejected. To do this, the justifier uses a few general types of question-answer plans. These typically require the justifier to trace backward along solution elements in the global data base from the questioned conclusion to the intermediate hypothesis or data that support it. Each step backward corresponds to the inference of one knowledge base rule. The justifier collects these intermediate inferences and translates them into English for presentation to the user.

There are a number of important distinctions between expert systems and conventional computer-based models which are highlighted in Diagram 1. The key distinctions are the inclusion of the knowledge base and a justifier in an expert system; both components are contained within the dashed rectangular area.

It was noted earlier that in an expert system the knowledge base and the control structure were kept separate. In a conventional model, the programmer's and domain specialist's expertise are embedded in the code. Thus, if one were to extract a dozen lines of code from a conventional model, it would be difficult to say what domain specific knowledge is being applied to the problem. In an expert system, a representative section of code may include a rule or provide information on the inferences being made. By maintaining a separation of the knowledge base and the control strategy, the knowledge in an expert system becomes more easily identified, more explicit, and more accessible. This implies that an expert system is both flexible and can be kept highly modular. Therefore, an expert system should be easier to change or augment than a conventional model. Since the domain knowledge is intermixed with the control structure (i.e., algorithms) in a conventional model, it is not as clear how the knowledge embedded in the code should change in order to improve the model.

The second key distinction between an expert system and a conventional model is the inclusion of a justifier. The justifier enables the user to make inquiries so that an "audit trail" from the problem's statement of its solution can be established. Expert systems are therefore more likely to be "transparent" to the user than conventional models. Furthermore, since the explanations provided by the expert system are in terms of its goals and its rules, they may be quite illuminating. The explanation capability may therefore facilitate the development of the system since it promotes a classification of previously unformalized knowledge in the problem domain.

Two other items which serve to distinguish an expert system from a conventional model are the elements marked with asterisks in Diagram 1. They are: (1) an agenda in the global data base; and (2) a scheduler in the control structure. In a conventional model, the plan for generating a solution is governed by one or more algorithms. The global data base thus consists of inputs, partial solutions and one or more actual solutions. Consequently, the agenda is predetermined by the plan (i.e., the algorithms employed); the plan thus becomes a part of the control structure. Because the rules are not explicit in a conventional model, the scheduler is also subsumed under the plan.

One of the primary reasons for the increased importance of expert systems within the modeling community is due to the tools which are available to the analyst. To better understand the capabilities provided by these tools, it is useful to discuss their role in the building of an expert system at two levels. The first level focuses on the traditional artificial intelligence programming languages. The second level focuses on what have become known as shells. A shell may be thought of as a specialty package which facilitates the construction of the knowledge base and provides the inferencing procedure for exploiting the information contained within the knowledge base.

If one were to build an expert system from scratch, one might choose an artificial intelligence programming language such as LISP or PROLOG. Within the United States, LISP (List Processor) is the artificial intelligence programming language of preference. Since LISP is a symbol manipulation language, it promotes the representation of knowledge which many authors claim to be oriented more towards symbolic processing than numeric processing. PROLOG (Programming in Logic) is widely used in Europe and Japan. The motivation behind PROLOG is to permit the analyst to specify the tasks in terms of logic statements rather than specify how they should be processed numerically by the computer. One difficulty with the use of artificial intelligence programming languages is the current lack of standardization. Consequently, many researchers have been moving toward the use of more universally accepted languages such as FORTRAN and PASCAL.

Since many applications share common methods for representing knowledge, performing inference, and maintaining large systems, researchers have designed shells which facilitate the development of expert systems. The motivation for developing a shell is to provide support to the analyst in representing symbolic knowledge. This knowledge can include facts, definitions, heuristic judgments, and procedures for doing a task or achieving a goal. The shell may be programmed in an artificial intelligence language, such as LISP in order to exploit its symbolic manipulation capabilities, or a common language such as FORTRAN. The main distinction between building an expert system based on an artificial intelligence or common language and a shell is that the analyst is relieved of the painstaking task of LISP or FORTRAN programming. Consequently, the shell enables the analyst (e.g., formulating rules or developing heuristics) and less on the production of code. The use of shells is therefore an important step in reducing the costs of developing an expert system.

The relationship between the idealized expert system and the various aspects of emergency management which follow will be illustrated through reference to the ORNL hazardous chemical spill program. This program has proved extremely useful to ORNL both for responding to emergency situations and as an educational tool.

Clearly, an accidental spill of oil or chemicals at ORNL may produce a serious emergency, depending on the properties and quantity of the substance released, the location, and whether or not the material enters a body of water. Emergency countermeasures must be applied immediately, particularly if the situation presents a health, safety, or environmental hazard. Appropriate countermeasures for various types of spill situations are governed by extensive reasoning. The U.S. National Oil and Hazardous Substances Pollution Contingency plan dictates the handling of many situations. Regional, Federal, State, city, and installation contingency plans also exist. In addition, at the individual installation level, a spill prevention, control, and countermeasures plan (SPCC plan) is required by Federal law for nontransportation facilities that may release oil spills to navigable waters.

As may be the case in other emergency response situations, the evaluation of the magnitude and severity of the spill are often carried out under hectic conditions before the scope of the spill emergency is understood. Mistakes in early containment, notifications, and hazard appraisal can be costly, and the more organized those early steps can be made, the better. Knowledge engineering may be viewed simply as a technique for formalizing common sense heuristic solutions into an understandable and computationally practicable form.

The program is potentially most useful during off-shift hours, when decisions must be made by personnel lacking formal training in the physical sciences. In such situations the program acts in a consulting mode while gathering information concerning a spill. Meanwhile, in the background, the program uses its various knowledge sources plus the accumulating facts to model and simulate the spill, to monitor for actions that should be taken or additional information that should be gathered, and to issue any appropriate messages and warnings.

As we shall see shortly, the ORNL hazardous chemical spill program includes some of the most important aspects of an emergency management problem. First, the system must integrate diverse sources of knowledge (e.g., SPCC plans, characteristics of each chemical, geographical features). Second, it must be capable of reasoning heuristically with incomplete and errorful data. Third, it must be able to accept data and advice continuously as they become available. Finally, it must be able to allocate limited resources to various tasks in a reasonable order.

II. The Nature of Emergency Management

The ability of public officials to respond correctly to an emergency situation is conditioned by their being able to recognize that a particular type of emergency has occurred and to initiate predetermined and ad hoc actions designed to alleviate the situation. Some localities have established contingency plans in the area of public safety. For example, New York City

has a computer-based system that directs the repositioning of fire trucks based on the location and severity of ongoing fires. Many metropolitan areas and local councils of governments have centralized the process by which an emergency is declared and resources positioned for action and response, e.g., during a hurricane or blizzard. Some localities have functioning command centers that enable them to pool, coordinate and allocate emergency resources. In general, however, the management of the response to an emergency tends to be reactive instead of proactive. Other aspects of emergency response are: (1) the need to determine appropriate guidelines for response actions and allocations of resources under changing conditions, (2) the training of emergency managers, and (3) the development of computer-based decision aids for use by emergency managers.

From a national perspective, FEMA has an important and direct role in ensuring that our citizens and their property are given the best protection under all emergencies. FEMA's major activity areas prior to an emergency are mitigation and preparedness. Given an emergency, FEMA is then concerned with response and recovery (see Chart 1).

For a major emergency, large quantities of information must be assembled from disparate sources, quickly processed and analyzed, and translated into a form useful for decisionmaking at many levels, up to and including the President. The collection and dissemination of such information is the function of the National Emergency Management System (NEMS). This system includes telecommunications, computers, and information on networks that link Federal and regional command centers. In terms of FEMA's coordinating and operational functions, our discussion is focused on the question "How can FEMA gather, analyze, and transmit information to all levels of emergency managers to assist them in determining the best response(s) to an emergency situation?" And more specifically, we are concerned with how expert systems can aid in answering this question.

III. The Framework for Emergency Mobilization¹

A. Background

1. National Security Decision Directive 47 sets forth principles and policies for emergency mobilization preparedness. It states that it is the policy of the United States to have an emergency mobilization capability that will ensure that government, at all levels, in partnership with the private sector and the American people, can respond decisively and effectively to any major national emergency.

¹Portions of Sections III and IV are abstracted from Federal Preparedness Circular 2, The Conceptual Framework for Emergency Mobilization Preparedness. It is important to point out that the discussion which follows is intended to be illustrative; it should not be viewed as a definitive statement of the activities involved in emergency mobilization preparedness.

FOUR PHASES OF COMPREHENSIVE EMERGENCY MANAGEMENT

1	MITIGATION (LONG-TERM) DEFINITION: Any activities which actually eliminate or reduce the probability of occurrence of a disaster. It also includes long-term activities which reduce the effects of unavoidable disasters.	PREPAREDNESS (TO RESPOND) DEFINITION: Preparedness activities are necessary to the extent that mitigation measures have not, or cannot, prevent disasters. In the preparedness phase, governments, organizations, and individuals develop plans to save lives and minimize disaster damage. Preparedness measures also seek to enhance disaster response operations.	RESPONSE (TO EMERGENCY) DEFINITION: Response activities occur immediately before, during and directly after an emergency or disaster. Generally, they are designed to provide warning, population protection, and emergency assistance. They also seek to reduce the probability of secondary damage and to speed recovery operations.	RECOVERY (SHORT AND LONG TERM) DEFINITION: Recovery continues until all systems return to normal or better. Short-term recovery returns vital life-support systems to minimum operating standards. Long-term recovery may continue for a number of years after a disaster. Their purpose is to return life to normal, or improved levels.
2	GENERAL MEASURES: Building codes Hazardous analyses Tax incentives/disincentives Zoning and land use management Bulldozer regulations/safety codes Compliance and enforcement Resource allocation/interstate sharing Preventive health care Public education Warning systems design/funding	Emergency Operations Plans Emergency exercises/training Warning systems installation Emergency communications systems Evacuation plans and training Resource inventories/contact lists Emergency Operations Centers Mutual aid agreements Public information/education Hazardous analysis updates	Activate plans, warnings, broadcast Notify public authorities Mobilize emergency teams/equipment Provide shelter, mass care Law & order, traffic control Emergency medical assistance Mass emergency operations centers Declare disaster/evacuate Search and rescue Emergency suspension of laws	Damage insurance/loans and grants Temporary housing Long-term medical care Disaster unemployment insurance Public information Health and safety education Reconstruction Counseling programs Economic impact studies Disaster Assistance Centers
3	HAZARD SPECIFIC MEASURES FLOOD: Dam construction/inspection Stream channel/levees/floodwalls Flood plain management Coastal zone management Flood insurance Contour farming/reforestation	Temporary levee construction Stream flow monitoring Ice and debris removal Temporary flood proofing Flood warning systems	Helicopter search Boat rescue Sandbagging Evacuation Damage assessment Mass shelter and care	Decontamination of water sources Clear debris and replant crops Pump out flooded basements Remove temporary floodproofing Monitor disease Repair and rebuild Vector control
	EPIDEMIC: School inoculations Rodent/insect eradication Water purification Sanitary waste disposal Health codes/laws/inspections Public health education	Stockpiling drugs Physician preparedness plans Public notification Quarantine regulations and plans Emergency medical authorities Inoculation programs, general	Quarantine Disinfect property Secure bodies Isolate carriers Public information Inoculation programs, specific	Continuing research into causes Long-term restorative therapy Case studies & evaluation Community health programs Infrastructural improvements
	FIRE: Fire codes No-smoking laws Fire zoning Fire safety education Sprinkler ordinances Smoke detector laws	Fire drills/exit signs Call boxes/smoke detectors Police crowd control training Fire department aid agreements Firefighter training Automatic sprinkler installation	Fire suppression Containment Crowd control Emergency medical services Mass shelter & care Investigation/assessment	Rebuilding Razing burned-out buildings Reforestation Incident evaluation Litigation Community reorganization
	HAZARDOUS MATERIALS SPILL: Transport speed limits Container structure codes Corporate licensing Restricted routing Materials identification codes	Containment and scrubbing equipment Stockpile neutralizing materials Emergency training for transporters Special apparatus for emergencies Responder training Chemical data information systems	Identify material Notify National Response Center/ CHEMTREC Containment, cordon area Plume tracking Air/water/soil contamination control	Reassess existing regulations Decontaminate environment Cleanup and disposal Public awareness Litigation Environmental impact statements
	LANDSLIDE: Forest management Preserve ground cover Maintain natural runoff Zoning & Land codes Real estate disclosure laws	Reinforce threatened structures Landslide monitors Stream rechannelization Erosion controls/ground covers	Assess stability of new formation Reinforce against secondary slip Evacuation Digout and debris removal	Reassess denuded areas New land-use planning Stabilize slopes Evaluate runoff, erosion Community awareness programs
	WIND: Roof anchors Window size and thickness codes Mobile home tie-downs Windbreaks	Storm shelter construction Property protection measures Storm watch and warning guides Tornado spotter network	Activate warning system Reinforce damaged property Broadcast all-clear Damage assessment	Reconstruction Debris clearance & removal Decontamination Water purification
	HURRICANE STORM SURGE: Barrier island zoning Coastal wetlands protection Replenish coastal sand dunes Construct breakwaters/levees Coastal zone management Public information programs	Vertical evacuation plans Storm tracking Shutter windows Seek shelter Evacuate planes and boats Tie-downs	Enforce evacuation Traffic routing & control Search & rescue Reinforce weakened structures	Rebuild destroyed sand dunes Bulkhead eroding areas Seed/stabilize eroding coast Evaluate/address secondary effects
	GASOLINE SHORTAGE: Alternatives research Allocations/international sharing Mass transit systems/car pooling Design energy efficient engines Reduce speed limits Energy conservation program	Stockpile reserves Reallocate to shortage areas Rationing plans Resource management plans	Odd-even purchase program Minimum purchase requirements Increase refinery production Deregulate oil Increase gasoline prices Activate alternatives Administer set-aside programs	Excess profits tax on companies Reassess allocation plan Two-way truck hauling Rebuild stockpiles Share resources
	NUCLEAR PLANT ACCIDENT: Site zoning Waste management/containment research Plant safety codes/inspections Plant operator training Environmental impact research/statements	Radiological emergency response plan Contamination monitoring Identify vulnerable populations Shelter preparation Designate Governor's technician Emergency procedures rehearsal Plant personnel training Community education program	Contain radioactivity Radiological monitoring/mapping Radiological exposure control Administer radio protective drugs Plume tracking Accident assessment Activate near-site Emerg. Opa. Ctr.	Reassess siting requirements Monitor deterioration of contaminant Decontamination Recovery and reentry On-site cleanup Investigation and reporting
	ATTACK: Continuity of government funding Coord. of defense/all-risk plans Maintain National Guard/Reserves Host area/hardened shelter dev't Diplomacy Arms control agreements	Defense mobilization plans Protection of vital records Continuity of government plans Nuclear Civil Protection plans Emergency resource mgmt plans Shelter management training War-related exercises	Disperse succession teams Deploy critical resources Activate rationing Activate vital workers plan Activate emergency broadcasts RADF monitoring	Redevelop/organize communities Reconstitute governments Provide security Allocate scarce resources

Often recovery includes mitigation measures against repeat of the disaster. Thus recovery and mitigation are related—showing circular nature of mitigation, preparedness, response and recovery.

EACH COLUMN SHOWS:

1. A definition of the phase;
2. Examples of general measures which apply to all hazards; and,
3. Examples of specific measures which apply to particular hazards.

This chart was developed for use by comprehensive emergency management review leaders. Many of the measures could be debated. We encourage clarification, editing and suggestions for making it more useful.
Contact: Hilary Whittaker, Director, Emergency Management Project
NATIONAL GOVERNORS' ASSOCIATION (202) 624-5365

NGA: 12/82

2. Executive Order 11490, Assignment of Emergency Preparedness Functions to Federal Departments and Agencies, requires departments and agencies to prepare national plans and programs and to attain an appropriate level of readiness with regard to the functions assigned. It also requires the Director of FEMA to establish Federal policies for and coordinate all emergency preparedness activities and functions of the Federal Government and be responsible for the preparation of guidance to Federal departments and agencies to assist them in performing their assigned emergency functions.

B. Definitions¹

1. Emergency. A sudden, generally unexpected event which does or could do harm to people, the environment, resources, property, or institutions. Emergencies range from relatively local events to regional and national events and may be caused by natural or technological factors, human actions, or national security-related events. The two types of emergencies are:
 - a. (1) National Security Emergency. Any natural, man-caused, or technological occurrence which, because of its size or intent, seriously degrades or threatens the national security of the United States.
 - b. (2) Domestic Emergency. A natural disaster or other emergency that does not seriously endanger national security.
2. Mobilization. The process of marshalling appropriate resources to manage emergencies. Mobilization as used here is defined broadly and applied to both civil and military resources.
3. Preparedness. A state of readiness to respond to and manage any type of emergency.

C. Dimensions of Emergency Mobilization

Three dimensions should be considered in emergency mobilization preparedness: the spectrum of emergencies, mobilization categories, and time phases.

1. Spectrum of Emergencies. Emergencies may be classified according to type and size. Consistent with the concept of integrated emergency management, plans for the full spectrum of emergencies, irrespective of type, should, to the extent possible, reflect common policies and compatible procedures so as to assure that efficient and effective use is made of all available capabilities.
 - a. Type of Emergency. Figure 1 shows examples of different emergencies, grouped according to type. Each type is then described.

¹All definitions are taken from Federal Preparedness Circular 2.

Figure 1. Examples of Types of Emergencies

NATURAL DISASTERS	TECHNOLOGICAL EMERGENCIES	RESOURCE SHORTAGES	DISORDERS	WAR
<p>Epidemics</p> <p>Extreme Hot or Cold Weather</p> <p>Fire</p> <ul style="list-style-type: none"> • Lightning • Spontaneous Combustion <p>Infestation</p> <p>Landshift</p> <ul style="list-style-type: none"> • Earthquake • Erosion • Mudslide • Volcanic Eruption <p>Snow & Ice</p> <p>Water</p> <ul style="list-style-type: none"> • Drought • Flood • Sea Surge • Tsunami <p>Wind</p> <ul style="list-style-type: none"> • Hurricane • Thunderstorm • Tornado 	<p>Dam Failure</p> <p>Explosion</p> <p>Fire</p> <p>Hazardous Materials Accident</p> <p>Pollution</p> <p>Radiological Accident</p>	<p>Energy & Fuel</p> <p>Food</p> <p>Manufactured Products</p> <p>Raw Materials</p> <p>Water</p>	<p>Arson</p> <p>Civil Disorder</p> <p>Crime</p> <p>Insurrection</p> <p>Labor Disturbance</p> <p>Riots</p> <p>Sabotage</p> <p>Terrorism</p>	<p>Biological or Chemical War</p> <p>Conventional War</p> <p>Nuclear War</p>

(1) Natural disasters occur as a result of nature and occasionally can be predicted; but, there is usually little that can be done to prevent them. Actions are needed to evacuate the affected population, take shelter, and/or mitigate the disaster's effects either before or after the event.

(2) Technological emergencies result from events which are primarily unintentional acts of man. Actions are needed to minimize occurrence and limit damage from those which do occur.

(3) Resource shortages occur when there is an insufficient supply of a resource on which society has become dependent. Actions are necessary to provide additional resources or substitute, when possible, and conserve existing supplies. Resource management programs should include standby plans and procedures for governmental intervention, as necessary, into the market system to ensure the enhancement of supply and the allocation of resources to military and essential civilian needs.

(4) Disorders are intentional acts of violence by one or more persons. They can sometimes be predicted and prevented. Proper action here is to emphasize prevention, be prepared to take immediate action to limit damage, and provide for recovery and restoration.

(5) War is an armed conflict between nation-states or between political factions within a nation-state. War is caused by man and can be prevented. Prevention (deterrence) is most desirable; but, if deterrence fails, U.S. objectives will be pursued by military or other means and will be guided by a plan for terminating the war in a manner that ensures achievement of those objectives.

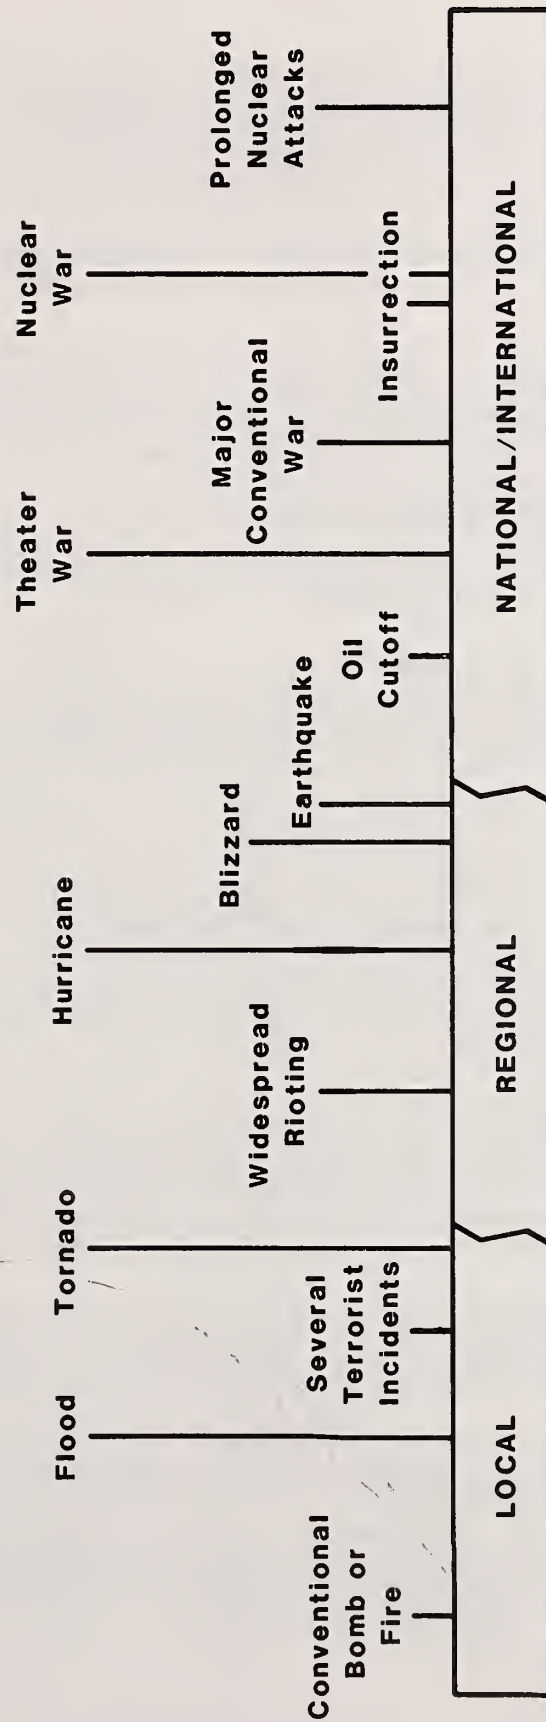
b. Size of Emergency. The various types of emergencies described above can be ranked by magnitude into local, regional, national, and international emergencies. It should be recognized, however, that there are variations in individual emergencies and that local emergencies may be of national or international interest and/or impact (see Figure 2).

c. Role of Government. Within the spectrum of emergencies, different levels of government play different roles in varied situations.

(1) In a domestic emergency, which occurs within a single geographic area, the situation can generally be managed by local and State resources. Federal assistance is provided only where authorized or required by law and when local and State governments need supplemental help.

(2) In a regional emergency which occurs across several States or within a single State, but with impact over several States, the Federal Government may coordinate the Federal response among several States and apply Federal resources across the entire region. Federal assistance is provided where authorized or required by law and to supplement, as necessary, local and State resources.

Figure 2. Spectrum of Emergencies



(3) In a domestic emergency, which impacts on the entire Nation, the Federal Government has the responsibility to coordinate the Federal response, provide assistance where authorized or required by law, and supplement local and State resources.

(4) In a conventional war, responsibility for management of national security emergencies rests primarily with the Federal Government.

(5) In a nuclear war situation, preparedness is a joint Federal, State, and local responsibility. Responsibility for national response and recovery is primarily a Federal responsibility dependent upon State and local coordinated support.

2. Mobilization Categories.

Seven mobilization categories are established according to the type of resource required to manage emergencies.

a. Military mobilization is the act of preparing for war or other emergencies through assembling and organizing military resources. It is the process by which the armed forces or part of them are brought to a state of readiness for war or other national emergency.

b. Industrial mobilization is the process of marshalling the industrial sector to produce goods and services, including construction, required to support military operations and the needs of the civil sector during domestic or national security emergencies.

c. Economic mobilization is the process of marshalling the money, credit, and taxes needed to: (1) finance the management of the emergency; (2) maintain a stable economy; and (3) stimulate key sectors of the economy.

d. Infrastructure mobilization is the process of marshalling the output of infrastructure systems to support the entire mobilization. Infrastructure systems include transportation, energy, communications, automated information processing, water, and agriculture.

e. Human resources mobilization is the process of marshalling people to provide needed labor. Human resources mobilization involves identifying and allocating human resources among competing demands.

f. Government mobilization is the process of marshalling resources of Federal, State, and local governments to carry out the tasks required to manage emergencies. It involves bringing to the appropriate state of readiness the leadership; policymaking groups; legislative bodies; courts; and supporting communications, facilities, procedures, and authorities to manage the emergency. Government mobilization activates and controls other aspects of mobilization.

g. Civil preparedness mobilization is the process of marshalling resources to provide protection for the people, industry, and institutions of the United States against the effects of the spectrum of emergencies. Civil mobilization involves providing for warning and emergency instructions to the public; relocation of people to safe areas; shelter, food, water, medical care, and other human needs; and recovery and reconstitution following the emergency.

3. Time Phases.

The final consideration in emergency mobilization preparedness is the timing of management actions. Four phases have been established for planning purposes. In a real emergency, the phases may not exist in clearly discernible stages, but they do provide a useful framework for planning for the full range of management actions. During each of these phases, the following types of actions are taken: mitigation, preparedness, response, and recovery.

Figure 3 shows the four mobilization phases.

a. Phase 1, Normal Operations Phase, is the base level state of activity which exists in the absence of an overt decision to take extraordinary measures. In this phase, mitigation actions may be directed to prevent or reduce the probability of occurrence of an emergency and minimize the potential for damage or loss of life.

b. Phase 2, Preparation Phase, involves taking preparedness actions to implement plans, procedures, and programs. It begins with a decision by appropriate authorities to increase the readiness of the Nation in one or more areas for a major domestic emergency or during periods of escalating international tensions.

c. Phase 3, Emergency Phase, involves a domestic or national security emergency, up to and including conventional or nuclear war. During this phase, response actions are taken directly before, during, and after the onset of the emergency to end the emergency, protect the public, provide critical assistance, limit damage, and reduce the probability of secondary effects.

d. Phase 4, Recovery Phase, involves restoring systems to normal. During this phase, short-term recovery actions are taken to assess damage and return vital life-support systems to minimum operating standards; long-term recovery actions may continue for many years.

D. Relationship Among the Spectrum of Emergencies, Mobilization Categories, and Time Phases.

Figures 4 and 5 show the relationship of the spectrum of emergencies, mobilization categories, and time phases.

Figure 3. Mobilization Phases




Figure 4. Conceptual Framework for Emergency Mobilization Preparedness

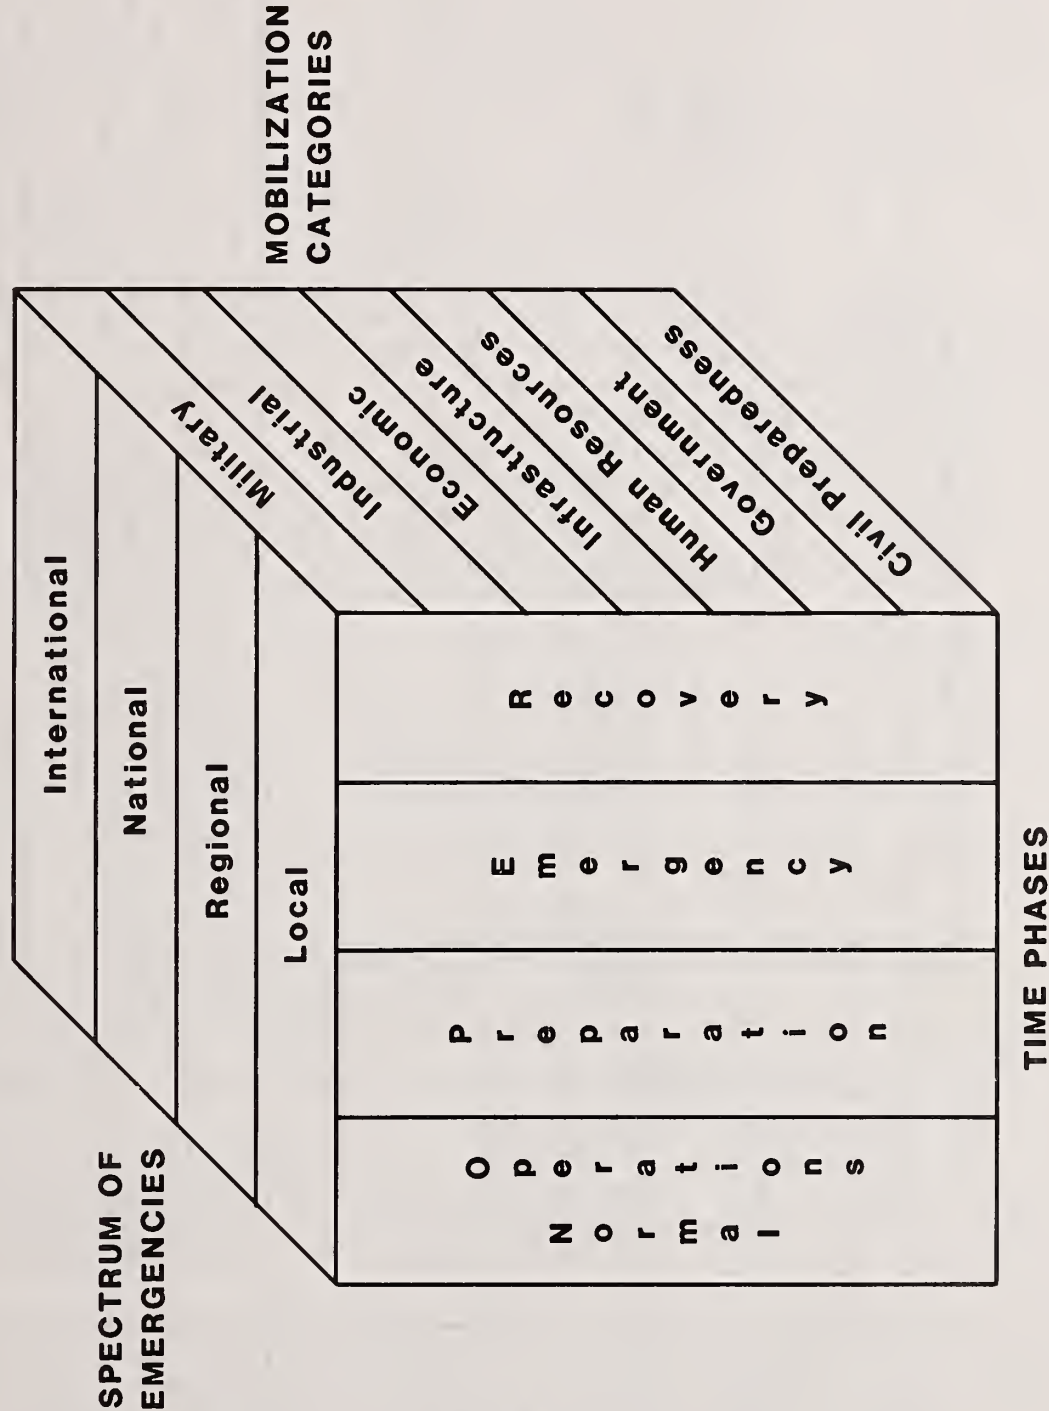


Figure 5. Time Phasing Across Mobilization Categories

TIME PHASES						
MOBILIZATION CATEGORIES	1 NORMAL OPERATIONS	2 PREPARATION	3 EMERGENCY PHASE			4 RECOVERY
			Domestic	National Security		
				Conventional War	Nuclear War	
MILITARY	Normal Readiness	Situation Monitoring	Selective Mobilization	Full or Total Mobilization	Full or Total Mobilization	Demobilization
INDUSTRIAL	Peacetime Production	Industrial Preparedness	Selective Surge	Defense Ind. Surge	Industrial Survival	Reconstitution of Industry
ECONOMIC	Normal Market. Normal Budgeting.	Normal Market. Accelerated Economic Activity. Planning Review.	Normal Market. Emergency Spending.	Special measures to expedite normal markets. Presidential discretion.	Extraordinary measures to expedite normal markets. Presidential discretion.	Phase out emergency measures
INFRASTRUCTURE	Peacetime Operations	Systems Preparedness	Selective Systems Surge	Wartime Systems Surge	Systems Survival	Reconstitution of Systems
HUMAN RESOURCES	Peacetime Services	Resource Assessment	Augmented Services	Wartime Services	Survival Services	Reconstitution of Services
GOVERNMENT	Normal Organization	Planning Review	Augmented Organization	Wartime Structure	Emergency Operations	Reconstitution of Government
CIVIL PREPAREDNESS	Program Planning	Increased Readiness	Response Operations	Response Operations	Relocation & Sheltering	Recovery Operations

IV. The Scope of Emergency Management Problems

In this section we describe the diversity of activities and related problems that FEMA must address in carrying out its mission. FEMA's activities include day-to-day data gathering, responding to the immediate needs of an emergency, and strategic planning and analysis. It is not our purpose to be complete here, but instead to highlight areas that may be addressed by advanced technological procedures, especially those of expert systems and artificial intelligence. Our aim is to emphasize the decisionmaking requirements faced by FEMA and related organizations and individuals.

Central to all of FEMA's activities is a viable and current information support system that is directed toward the needs of emergency managers. Such a system is defined to include not only data and descriptive information but also analysis and interpretive programs. The establishment and maintenance of this support system, although a problem in itself, cannot be accomplished without recognition of the decision problems and other related problems of emergency managers. These include the development and implementation of training procedures for emergency managers at all levels, the requirement of emergency managers at all levels to be able to find and interpret current emergency response rules and regulations, the need of FEMA to be able to evaluate proposed emergency guidelines and legislation, and the basic management procedural requirements for coordinating multi-regional responses and allocation of resources. Thus, to be effective, an information support system must be designed to reflect the needs of these varied problems. In particular, we must be cognizant of the many confounding and limiting aspects inherent in the broad spectrum of emergency operations. These include the varied nature of emergencies and the extremely wide range of possible responses requiring formal evaluation procedures that, as of today, do not exist. At best, such procedures are now based on past experiences, which are not necessarily generalizable, and current, untested doctrines. The operational experiences of field emergency managers are not extensive, and, as there is much turnover in these positions, the "hands-on" knowledge that exists in the field fluctuates and is never at a high level of expertise. Compounding an emergency operation decision process is the lack of complete information, and whatever information is available is often conflicting. But, decisions must be made. By its very nature, an emergency is a complex situation that often requires a nonstandard response: a response that has been evaluated in terms of the available information. The capability of emergency managers to assimilate and evaluate changing situations is a function of their ability for analyzing available information and their past experiences. And, as we have noted, these capabilities are rather limited.

The time taken to respond to an emergency can be critical. Public safety organizations, e.g., police, fire, ambulance service, are always on duty and take pride in their short response times. But for broader emergency situations, like toxic and explosive spills or hurricanes that require decision structures and response systems that are basically dormant, the ability to mount a proper and timely response is difficult. Here, emergency managers must be able to couple their static personnel and resource base with sporadic incoming information to make a dynamic set of responses. This can only be effective if a full range of contingency plans has been developed and tested by the emergency management team. Such contingency planning is not a

priority item at most levels of emergency management. It is costly; it requires experienced personnel; but, it is a tactical activity that must be placed on a solid base.

The analysis of emergencies and emergency responses is backed up with some theory on the physics of particular emergencies, but little theory on the behavioral aspects of most emergencies. There are fire spread models, nuclear blast and fallout models, network flow models, epidemiological models, some of which can be of assistance in contingency planning, as well as in an operational situation. A key question is how available theory can be utilized in an emergency information support system, especially at the local level. One is quick to recognize that the major concerns of emergency management operations, like all complex decision situations, have both technological and human components. And, as the problem is of national importance and something must be done to understand better all aspects of emergency management, there is a danger of rushing to apply technical methodologies. An integrated emergency information-communications system will be ineffective if we ignore the personnel and behavioral aspects of who is to use it, how will the users be trained, and what the decision and resource needs of the users are. To propose an expert system is one thing, but to develop one that can be shown to be of value to emergency managers is quite another matter. Cost-benefit analyses of such proposals are difficult; but performing such an analysis for a proposed expert system will highlight the strong and weak aspects and, thus, make for better choices. The full range of emergency management operations must be laid out and subjected to analysis in terms of needs and priorities.

V. The Managers of Emergencies

From a general emergency perspective, the individuals and groups that concern themselves with an emergency, i.e., the stakeholders and actors, include persons at the Federal, State and local levels, as well as persons in private sector organizations. Thus, emergency management operations and supporting systems must be responsive to a wide diversity of needs and priorities. A listing of Federal groups includes FEMA officials, emergency managers and emergency response systems; The White House; Congress; Departments of Transportation, Energy, Commerce, and Interior; Environmental Protection Agency; Nuclear Regulatory Commission; and the National Oceanographic and Atmospheric Administration, among others. Of course, the interactions and needs of these groups change by type of emergency. Depending upon the stage of an emergency, most information is used for status reports and not for decisionmaking. However, as an emergency is not static, the needs of the Federal stakeholders and actors can change rapidly from a passive to an active mode. Thus, the support requirements are quite complex, and, at this point in time, not well understood.

At the State level, the activities of emergency management are usually centralized within the governor's office and supervised by the State Director of Emergency Management. As is often the case, emergencies know no State lines or regional boundaries. Thus, coordination efforts must be developed among multi-state and multi-regional agencies, as well as between individual State agencies. Here, resources such as the National Guard and the full range of State safety personnel have to be allocated and coordinated based upon the changing demands of an emergency. Officials, public safety officers, and

volunteers on the firing line at the local level have similar requirements. Coordination is a must for the impacted city and county agencies and for local councils of government. Public safety at this lower level requires fire, police, and other emergency services to be responsive to the decisions of the respective chiefs or directors of public safety. They must be ready to act on decisions made at the scene and on the decisions of those evaluating the overall needs of the threatened area. Decisions on when and how to evacuate an area must not be made in an ad hoc manner, but be based on tested contingency plans and the most current evaluation of the emergency conditions. The integration of private sector groups such as the Red Cross, the National Voluntary Organizations Active in Disasters (NVOAD), and the Civil Defense National Radio System must be coordinated at all levels. How best to manage these diverse groups and how best to manage the data, information, and decision needs of the stakeholders and actors is the major research task facing FEMA; and, it, is the area in which we see expert systems being able to make a major contribution.

VI. The Range of Emergencies

In this section, we offer without comment a listing of emergencies that are of concern to FEMA and/or State and local agencies:

- Nuclear plant: meltdown, radiation fallout
- Earthquake: dams, spills, fire, landslides
- Hurricane: winds, flooding, tornadoes
- Floods
- Volcano eruption
- Mobilization
- Enemy attack
- Civil disorder
- Terrorist attack
- Epidemic
- Power failure
- Resource shortage
- Economic emergency
- Agricultural emergency
- Boating, ships, ferry accidents
- Airplane crash
- Fire: city, forest, chemical, gas, oil
- Explosion: chemical, gas, oil
- Mine disaster
- Hazardous spills
- Toxic airborne release
- Subway/Transportation disaster
- Dam Failure
- Water shortage/pollution
- Air pollution

Fortunately, we in the United States encounter a small subset of such emergencies and those that we do encounter are usually at a low level of intensity. However, our Federal emergency disaster costs are quite high. For example, as reported by the United Press International, the United States spent \$552 million on disaster relief in 1984. During 1984 the President directed FEMA to allocate supplemental Federal aid 38 times, the greatest number of

allocations since 1979 and 11 more than the annual average over the past decade. Twenty-six of the 38 disasters were due to floods that affected 33,800 families. A total of 41,000 families sought Federal aid. The most costly disaster (\$164 million) occurred in New York and New Jersey and was due to a severe storm and torrential rains. Other disasters included tornadoes, citrus freeze, and fires.

VIII. The Management Problem

The full range of the management problem is described in the terms of Chart 1: Four Phases of Comprehensive Emergency Management. We briefly describe aspects of these phases and associated activities.

A. Mitigation

The basic activities of mitigation, which are carried out at all levels of government are legislation, codes, prevention, and education. Decisionmaking here tends to be concentrated in setting priorities and the allocation of funds and other resources. Computer-based aids include information retrieval systems for analyzing current codes and legislation and econometric models for analyzing the cost-benefit aspects of proposed regulations.

B. Preparedness

The activities in this phase involve many of the stakeholders and actors across a wide range of activities. Here we are concerned with training, resource positioning, status reporting and communications, shelter and medical facilities, evacuation procedures, contingency planning, and education. Decisions need to be made that relate to the succeeding phases of response and recovery. Computer-based aids include location models, man-machine simulation exercises, routing procedures, and automated detection systems (e.g., for earthquakes). Expert systems could have an important role in training and contingency planning.

C. Response

The full range of stakeholders and actors are involved in this phase. The main activities are information gathering and analysis, allocation of resources, coordination of responses and resources, evaluation of proposed responses, evacuation, and protection. The main technological aids are communications and information processing. It is this phase that requires a wide range of expertise, e.g., transportation, traffic control, search, handling of toxic materials, nuclear plant operations, radiation and plume dispersion, weather, medical, and public safety. It is during this phase of an emergency that expert systems have the greatest potential.

D. Recovery

This last phase of an emergency includes the activities of insurance and loans, cleanup and reconstruction, assessment, and post audit of the emergency response system. Federal and local authorities can be assisted in this phase by expert systems and other data processing aids that enable them to determine if a request for assistance can be given within the restrictions of the applicable legislation and regulations. Post-audit of an emergency is important so as to be able to feed back to the response system information on what worked and what did not. Such feed back is important to an expert system and represents the true nature of artificial intelligence.

APPENDIX II: EXPERT SYSTEMS AND EMERGENCY MANAGEMENT: RESEARCH DIRECTIONS

Saul I. Gass
College of Business and Management
University of Maryland
College Park, MD 20742

and

Robert E. Chapman
Mathematical Analysis Division
Center for Applied Mathematics
National Bureau of Standards
Gaithersburg, MD 20899

This appendix reproduces material that describes possible research directions in expert systems and emergency management for FEMA. It is based on the presentations at the Symposium and a follow-up Workshop on Expert Systems Applications and FEMA Research Activities held on April 26, 1985 at FEMA Headquarters.

EXPERT SYSTEMS AND EMERGENCY MANAGEMENT: RESEARCH DIRECTIONS

This report is based on presentations and discussions of the First Symposium on the Theory and Application of Expert Systems in Emergency Management Operations (April 24-25, 1985) and the related Workshop on Expert Systems Applications and FEMA Research Activities (April 26, 1985). The purpose of this report is to describe research directions in expert systems for consideration by FEMA management.

A general conclusion is that the new and emerging field of expert systems could be of value to Federal and local emergency managers. Based on the materials presented at the Symposium and Workshop, we have evidence that expert systems (ES) have already been integrated with emergency management (EM) systems and have proven to be worthwhile decision-aids. However, as the full realization and benefits of expert systems, as well as its parent field of artificial intelligence (AI), are yet to be established, there is a need for FEMA to proceed with caution in its commitment of funds and reliance on a future payoff.

FEMA's role in an ES/EM program can take on many directions:

- projects and research activities that support FEMA Headquarters (operational activity)
- prototype projects designed to support State and Local projects (demonstration activity)
- establishing a central source for dissemination of ES/EM publications (clearinghouse activity)
- holding of symposia and workshops on state-of-the-art ES/EM developments (knowledge transfer activity)
- setting standards for and the dissemination of ES/EM software and related computer-based training aids that support the operations of local emergency centers (technology transfer activity)
- developing ES/EM training exercise to be used at the National Emergency Training Center and at the local level (education activity)

To date, in association with the National Bureau of Standards, and the Department of Transportation, FEMA is supporting ES/EM applications in data error analysis (NBS), message management systems (NBS), and transportation network analysis (DOT). These projects are designed to support elements of FEMA Headquarters. FEMA has also purchased the software package TIMM (The Intelligent Machine Model), an aid for the development of expert systems, that

will be used in support of the message management and transportation projects. The question to be addressed now is "What other ES/EM areas should FEMA support?" Our response to this question is in terms of possible demonstrations and research activities. First, we give an unordered list of emergencies and corresponding areas that appear to be within the domain of ES/EM (see Figure 1). The items shown in Figure 1 are directed at the response phase of emergencies. It is in this phase that expert systems appear to have the greatest potential. Note that most of the response-oriented expert systems would be operated by State and local agencies. At the National level, FEMA's response-oriented expert systems could be applied to issues such as mobilization and enemy attack.

Prototypical FEMA Expert System Activities

The areas of expert systems applications shown in Figure 1 are the following:

1. The problems of moving people, material and equipment into and out of an emergency zone (evacuation routing)
2. The location and assignment of evacuees to emergency shelters (shelter allocation)
3. The determination that an emergency situation is likely to occur (early warning)
4. The systemic and decision processes by which a determination is made that an emergency is occurring and automatic controls are to be initiated to shut down or degrade essential equipment and services (automated response/shutdown)
5. The problems associated with the marshaling of the full range of resources necessary for the proper response to an emergency (resource management)
6. The analysis and determination of the source, type and extent of certain emergency situations (fault analysis)

We next discuss items (1), (2) and (5) in terms of possible FEMA activities and research in expert systems. As noted below, expert system requirements for these areas are interrelated with shelter allocation dependent on routing decisions, and resource management dependent on both shelter allocation and routing decisions.

The problem of evacuation routing is critical in those emergencies that involve the general public to a high degree. Recent chemical spills and toxic airborne releases have captured the headlines and have raised serious questions as to whether communities that live in the shadow of chemical plants are in a position to respond to a serious event. Thus, there is a need to establish and test contingency plans for notifying the surrounding community that an emergency has occurred, with the initiation of specific plans based on an expert system's analysis of the emergency, weather, time of day and other critical factors. The development of a prototype expert system for evacuation

Figure 1. Emergency/Expert Systems Areas of Application

EMERGENCY	EXPERT SYSTEM AREA OF APPLICATION					
	Evacuation routing	Shelter allocation	Automated response/ shut down	Resource management (personnel, equipment, communications, information processing)	Fault (source) analysis	
Nuclear plant disaster	•		•		•	
Earthquake	•		•	•		
Hurricane	•		•	•		
Floods	•	•	•	•		
Volcanoes	•	•	•	•		
Enemy attack	•		•	•		
Civil disorder			•	•		
Terrorist attack			•	•		
Epidemic	•		•	•	•	
Power failure			•	•	•	
Resource shortage			•	•		
Economic emergency			•	•		
Agricultural emergency			•	•		
Boating, ships, ferry accidents			•	•		
Airplane crash			•	•	•	
Fire	•	•	•	•	•	
Explosion	•		•	•	•	
Mine disaster	•		•	•	•	
Hazardous spills	•		•	•	•	
Toxic airborne release	•	•	•	•	•	
Transportation disaster	•	•	•	•	•	
Dam failure	•	•	•	•	•	
Water shortage/pollution	•	•	•	•	•	
Air pollution	•	•	•	•	•	

routing, designed and tested for a specific community, should be investigated by FEMA. The current DOT transportation expert system, plus the extensive literature in routing procedures, would be of value here.

For those emergencies in which evacuation is required, shelter allocation must be coordinated with routing assignments. Hence, a more general expert system that integrates the emergency management functions of escape routing and shelter allocation would be of value. The requirements of a prototype expert system that includes both routing and allocation should be studied for a high-risk community.

As evidenced by Figure 1, the problem of general resource management (personnel, equipment, communications, information processing) is one that cuts across all emergencies; it represents a fruitful area for FEMA to investigate. As a generic expert system for resource management that is applicable to the full range of emergencies is not possible, attention should be focused on resource problems associated with evacuation type problems (e.g., toxic release, hurricane). For these problems, we feel that the expert system requirements of resource management and those of routing and shelter allocation can be combined into a powerful emergency control tool. Further, such an integrated system, when coupled with local and regional emergency information and data systems, would be able to assess changing conditions and thus, be an extremely valuable decision aid to local emergency managers. Specifications for such an expert system based on the requirements of a specific community should be developed.

Within the other phases of emergency management - mitigation, preparedness and recovery - expert systems can also be of value. For example, expert systems can play an important role in preparedness if they can be integrated as decision aids into training exercises, with the expert systems being those that the trainee would actually use in the field. During a recovery phase, expert systems can be used to analyze eligibility for loans and government reconstruction programs.

With the expected heavy use of personal computers by field personnel, FEMA should investigate which field-based expert system activities can be generalized and made available.

FEMA Expert Systems Projects

This section describes expert system applications in areas that stem from FEMA's operational and educational activities; these areas are of broader scope and/or are more technically oriented than the prototypical ones described above. In particular, we address how expert systems can be used by FEMA in the analysis of computer-based model outputs, automated early warning, training and research, and information retrieval; what areas of expert system research would be of most value to FEMA; and conclude with a short list of general expert system activities that should be considered by FEMA.

1. Model Output Reviewer

FEMA uses many large-scale econometric, input-output and other models to analyze impacts of emergencies on the nation's industrial and transportation systems. Such analyses are often based on changing scenarios and require the review of a great deal of computer output. It has been proposed that FEMA investigate a computer-based process by which such outputs can be reviewed automatically, and the essence of the data be processed and formatted by an expert system that replicates what the human expert analysts would accomplish. At a minimum, such an expert system would be able to organize the model outputs into tables and graphs; at a maximum, it would be able to find, filter and adjust conflicting data and to make recommendations for further analyses and specific actions. Output analysis programs have been used for quite some time, but here we are proposing a major advance in such post-processor programs. This project should be able to use the expert system data analysis system being developed by NBS for FEMA; also, the construction of the expert system model output reviewer should be facilitated by the use of TIMM. Issues that would have to be resolved include the specification of the models, the associated expert(s) to be used to build a knowledge base, and the types of modeling exercises to be reviewed.

2. Automated Early Warning/Response

Some emergencies such as earthquakes or nuclear power-plant accidents require a means for early detection and, given that there is no means of stopping the incident, a quick response e.g., within (100 seconds) is required to shutdown or softly degrade critical systems and to warn emergency managers. It has been suggested that an expert system, tied into a detection network, would be able to automatically shutdown critical parts of the infrastructure and issue the appropriate warnings. The expert system would be designed to interpret the information produced by the detection network and to make the "call" that a destructive incident is about to occur. There are a number of issues here that go beyond just being able to build such an expert system. For example, our present ability to differentiate earthquake signals is not perfected and false alarms would trigger needless actions (although the expert system triggering mechanism could be designed not be activated unless an earthquake of sufficient Richter magnitude was actually happening). Furthermore, critical facilities may not consent to be tied into such an automatic system. From an expert system point of view, the logic design would need a proper research base and an ability to be tested and validated. Personnel at the Institute of Safety and Systems Management, University of Southern California have undertaken some research in this area.

3. Expert System for Training and Research

It has been proposed to investigate how expert system concepts can be integrated within the present training program of the FEMA National Emergency Training Center (NETC). The current Central City Training exercise could be augmented by expert system decision aids for evacuation routing, allocation of resources, and shelter allocation. Using the NETC would enable FEMA to test out the utility of expert systems, and if successful, initiate development of operational systems for transfer to field locations. Research would be required to scope the application areas, and to design the decision rules and

algorithmic processes for evaluation of the exercise information. This concept can also be extended to FEMA Headquarters' exercises in which expert systems can be used as decision aids and for automating certain exercise functions (e.g., the message management system being developed by NBS).

4. Information/Knowledge Base Access

Any emergency, especially at the national level, requires access to a variety of data bases and, given expert systems will be used in the emergency management area, access to many types of knowledge bases. The proper use of such information can be a complex matter and could possibly require an expert system, or certainly an advanced information retrieval system. At issue here is the structuring and availability of the data and knowledge bases, the usages and analyses to be made of the information, and the availability of the proper development tools.

5. Related Research Proposals

As noted, the expert system/artificial intelligence field has not reached its full potential and is under constant development. A number of ES/AI research topics that would hopefully be of benefit to the emergency management field have been proposed. We cite them here, although we feel that, in general, due to resource constraints and uncertainty in the ES/AI field, FEMA should not be in the forefront of advanced research, but FEMA's expert system activities should be directed towards utilization of tested expert system theory and practice.

a. Natural Language Dialogue Interface

This proposal is to improve the ability to use natural English and the emergency management jargon for those who interface with computer and related information systems. The hope is to ease the use of computers by non-technical emergency management personnel by enabling them to hold a dialogue with the computer system to retrieve information and initiate analyses.

b. Decision Rules in Expert Systems

There are a number of basic areas of interest here such as:

- (1) The manner in which an ES analyzes uncertainty to determine certainty factors.
- (2) The ability to state a full set of contingencies and decision rules.
- (3) The way partial and imprecise information is interpreted.
- (4) The means by which an expert system "learns" during its lifetime.
- (5) The use of decision models such as the Analytic Hierarchy Process within an expert system.

c. Validation of ES

As with all decision models, there is a need to validate the recommendations of an expert system in terms of their consistency and acceptability by emergency managers, and their ability to replicate expert decisions. The process of validation is extremely important with respect to FEMA expert systems, as the automation of decisions that impact human lives is always open to second guessing and the clarity of hindsight. Thus, for an ES, FEMA must establish a process with which the ES is validated by the associated experts and field-tested. The design of validation processes and the conducting of associated experiments are expert system aspects of prime importance to FEMA.

d. Expert Systems Shells and PC Expert Systems

The building of an expert system can be facilitated greatly by using commercially available software and/or specialized equipment. Expert system shells (i.e., software that can handle the logic and construction of decision rules and is independent of the applications) need to be evaluated in terms of FEMA requirements. In addition, with the expected heavy use of personal computers (PC's) by field personnel, FEMA should evaluate commercial PC/ES based on field requirements, as well as FEMA Headquarters' applications.

6. General FEMA Activities

- a. Standards and guidelines for the development of expert systems by FEMA and field personnel.
- b. Expert system technology clearinghouse for distribution of emergency management related systems and information.
- c. Develop expert system curriculum and materials for training of emergency management personnel.
- d. Investigate ES/EM activities being done by other countries.
- e. Tracking of other fields use of ES/AI and microcomputers and relationship to emergency management.
- f. Setting FEMA expert system priorities.

U.S. DEPT. OF COMM. BIBLIOGRAPHIC DATA SHEET (See instructions)		1. PUBLICATION OR REPORT NO. NBS/SP-717	2. Performing Organ. Report No.	3. Publication Date November 1986
4. TITLE AND SUBTITLE Theory and Application of Expert Systems in Emergency Management Operations (Proceedings of a Symposium)				
5. AUTHOR(S) Saul I. Gass, Robert E. Chapman, editors				
6. PERFORMING ORGANIZATION (If joint or other than NBS, see instructions) NATIONAL BUREAU OF STANDARDS Federal Emergency Management DEPARTMENT OF COMMERCE Agency, Washington, DC 20472 Gaithersburg, MD 20899			7. Contract/Grant No.	8. Type of Report & Period Covered Final
9. SPONSORING ORGANIZATION NAME AND COMPLETE ADDRESS (Street, City, State, ZIP) National Bureau of Standards Federal Emergency Management Agency Department of Commerce Washington, DC 20472 Gaithersburg, MD 20899				
10. SUPPLEMENTARY NOTES Library of Congress Catalog Card Number 86-600546 <input type="checkbox"/> Document describes a computer program; SF-185, FIPS Software Summary, is attached.				
11. ABSTRACT (A 200-word or less factual summary of most significant information. If document includes a significant bibliography or literature survey, mention it here) The First Symposium on The Theory and Application of Expert Systems in Emergency Management, held at the Department of Commerce, Washington, D.C. (April 24 and 25, 1985) was funded by the Federal Emergency Management Agency and organized by the National Bureau of Standards' Operations Research Division. The purpose of the symposium was to bring together researchers in expert systems, artificial intelligence, and emergency operations in a forum to review the concepts of expert systems and the problems of emergency management, with the objective of determining how expert systems can be used to augment the experience of local, State and Federal emergency managers faced with the difficult tasks of determining the best response to an emergency situation. Speakers addressed the following areas: The theory and uncertainty aspects of expert systems, artificial intelligence's future role in emergency management, technology for building and using expert systems, emergency management decisions and information needs and uses, applications of expert systems in the management of chemical spills and shipboard and coal mine fires, and the role and use of simulation in emergency management expert systems.				
12. KEY WORDS (Six to twelve entries; alphabetical order; capitalize only proper names; and separate key words by semicolons) artificial intelligence; chemical spills; coal mine fires; decision support systems; emergency management; expert systems; fire management; shipboard fires; simulation; uncertainty				
13. AVAILABILITY <input checked="" type="checkbox"/> Unlimited <input type="checkbox"/> For Official Distribution. Do Not Release to NTIS <input checked="" type="checkbox"/> Order From Superintendent of Documents, U.S. Government Printing Office, Washington, D.C. 20402. <input type="checkbox"/> Order From National Technical Information Service (NTIS), Springfield, VA. 22161			14. NO. OF PRINTED PAGES 264 15. Price	

NBS *Technical Publications*

Periodical

Journal of Research—The Journal of Research of the National Bureau of Standards reports NBS research and development in those disciplines of the physical and engineering sciences in which the Bureau is active. These include physics, chemistry, engineering, mathematics, and computer sciences. Papers cover a broad range of subjects, with major emphasis on measurement methodology and the basic technology underlying standardization. Also included from time to time are survey articles on topics closely related to the Bureau's technical and scientific programs. Issued six times a year.

Nonperiodicals

Monographs—Major contributions to the technical literature on various subjects related to the Bureau's scientific and technical activities.

Handbooks—Recommended codes of engineering and industrial practice (including safety codes) developed in cooperation with interested industries, professional organizations, and regulatory bodies.

Special Publications—Include proceedings of conferences sponsored by NBS, NBS annual reports, and other special publications appropriate to this grouping such as wall charts, pocket cards, and bibliographies.

Applied Mathematics Series—Mathematical tables, manuals, and studies of special interest to physicists, engineers, chemists, biologists, mathematicians, computer programmers, and others engaged in scientific and technical work.

National Standard Reference Data Series—Provides quantitative data on the physical and chemical properties of materials, compiled from the world's literature and critically evaluated. Developed under a worldwide program coordinated by NBS under the authority of the National Standard Data Act (Public Law 90-396).

NOTE: The Journal of Physical and Chemical Reference Data (JPCRD) is published quarterly for NBS by the American Chemical Society (ACS) and the American Institute of Physics (AIP). Subscriptions, reprints, and supplements are available from ACS, 1155 Sixteenth St., NW, Washington, DC 20056.

Building Science Series—Disseminates technical information developed at the Bureau on building materials, components, systems, and whole structures. The series presents research results, test methods, and performance criteria related to the structural and environmental functions and the durability and safety characteristics of building elements and systems.

Technical Notes—Studies or reports which are complete in themselves but restrictive in their treatment of a subject. Analogous to monographs but not so comprehensive in scope or definitive in treatment of the subject area. Often serve as a vehicle for final reports of work performed at NBS under the sponsorship of other government agencies.

Voluntary Product Standards—Developed under procedures published by the Department of Commerce in Part 10, Title 15, of the Code of Federal Regulations. The standards establish nationally recognized requirements for products, and provide all concerned interests with a basis for common understanding of the characteristics of the products. NBS administers this program as a supplement to the activities of the private sector standardizing organizations.

Consumer Information Series—Practical information, based on NBS research and experience, covering areas of interest to the consumer. Easily understandable language and illustrations provide useful background knowledge for shopping in today's technological marketplace.

Order the **above** NBS publications from: Superintendent of Documents, Government Printing Office, Washington, DC 20402.

Order the **following** NBS publications—FIPS and NBSIR's—from the National Technical Information Service, Springfield, VA 22161.

Federal Information Processing Standards Publications (FIPS PUB)—Publications in this series collectively constitute the Federal Information Processing Standards Register. The Register serves as the official source of information in the Federal Government regarding standards issued by NBS pursuant to the Federal Property and Administrative Services Act of 1949 as amended, Public Law 89-306 (79 Stat. 1127), and as implemented by Executive Order 11717 (38 FR 12315, dated May 11, 1973) and Part 6 of Title 15 CFR (Code of Federal Regulations).

NBS Interagency Reports (NBSIR)—A special series of interim or final reports on work performed by NBS for outside sponsors (both government and non-government). In general, initial distribution is handled by the sponsor; public distribution is by the National Technical Information Service, Springfield, VA 22161, in paper copy or microfiche form.

U.S. Department of Commerce
National Bureau of Standards
Gaithersburg, MD 20899

Official Business
Penalty for Private Use \$300